NBER WORKING PAPER SERIES

IMPERFECT COMPETITION, COMPENSATING DIFFERENTIALS AND RENT SHARING IN THE U.S. LABOR MARKET

Thibaut Lamadon Magne Mogstad Bradley Setzler

Working Paper 25954 http://www.nber.org/papers/w25954

NATIONAL BUREAU OF ECONOMIC RESEARCH 1050 Massachusetts Avenue Cambridge, MA 02138 June 2019

The opinions expressed in this paper are those of the authors alone and do not reflect the views of the Internal Revenue Service or the U.S. Treasury Department. This work is a component of a larger project on income risk in the United States, conducted through the SOI Joint Statistical Research Program. We are grateful to Raj Chetty, Nathan Hendren, Danny Yagan and Owen Zidar for help and guidance in using the IRS data. We also appreciate the constructive comments and suggestions from Thomas Lemieux, three anonymous referees, Neele Balke, Stephane Bonhomme, Derek Neal, and discussants and participants at various conferences and seminars. Mogstad and Setzler acknowledge funding from the Washington Center for Equitable Growth and the Norwegian Research Council. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Thibaut Lamadon, Magne Mogstad, and Bradley Setzler. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Imperfect Competition, Compensating Differentials and Rent Sharing in the U.S. Labor Market Thibaut Lamadon, Magne Mogstad, and Bradley Setzler NBER Working Paper No. 25954 June 2019 JEL No. J20,J30,J42

ABSTRACT

The primary goal of our paper is to quantify the importance of imperfect competition in the U.S. labor market by estimating the size of rents earned by American firms and workers from ongoing employment relationships. To this end, we construct a matched employer-employee panel data set by combining the universe of U.S. business and worker tax records for the period 2001-2015. We describe several important features of the U.S. labor market, including the size of firm-specific wage premiums, the sorting of workers to firms, the production complementarities between high ability workers and productive firms, and the pass-through of firm and market shocks to workers' wages. Guided by these empirical results, we develop, identify and estimate an equilibrium model of the labor market with two-sided heterogeneity where workers view firms as imperfect substitutes because of heterogeneous preferences over non-wage job characteristics. The model allows us to draw inference about imperfect competition, compensating differentials and rent sharing. We also use the model to quantify the relevance of non-wage job characteristics and imperfect competition for inequality and tax policy, to assess the economic determinants of worker sorting, and to offer a unifying explanation of key empirical features of the U.S. labor market.

Thibaut Lamadon Department of Economics University of Chicago 1126 East 59th Street Chicago, IL 60637 and NBER lamadon@uchicago.edu

Magne Mogstad Department of Economics University of Chicago 1126 East 59th Street Chicago, IL 60637 and NBER magne.mogstad@gmail.com Bradley Setzler Department of Economics University of Chicago 1126 East 59th Street Chicago, IL 60637 bradley.setzler@gmail.com

A data appendix is available at http://www.nber.org/data-appendix/w25954

1 Introduction

How pervasive is imperfect competition in the labor market? Arguably, this question is really about the size of rents earned by employers and workers from ongoing employment relationships (Manning, 2011). In the textbook model of a competitive labor market, the law of one price holds and there should exist a single market compensation for a given quality of a worker, no matter which employer she works for. If labor markets are imperfectly competitive, however, the employer or worker or both may also earn rents from an existing employment relationship. If a worker gets rents, the loss of the current job makes the worker worse off—an identical job cannot be found at zero cost. If an employer gets rents, the employer will be worse off if a worker leaves – the marginal product is above the wage and worker replacement is costly.

To draw inference about imperfect competition in the labor market, it therefore seems natural to measure the size of rents earned by employers and workers. However, these rents are not directly observed, and recovering them from data has proven difficult for several reasons.¹ One challenge is that observationally equivalent workers could be paid differently because of unobserved skill differences, not imperfect competition (see e.g. Abowd et al., 1999; Gibbons et al., 2005). Another challenge is that observed wages may not necessarily reflect the full compensation that individuals receive from working in a given firm. Indeed, both survey data (e.g., Hamermesh, 1999; Pierce, 2001; Maestas et al., 2018) and experimental studies (e.g., Mas and Pallais, 2017; Wiswall and Zafar, 2017) suggest that workers may be willing to sacrifice higher wages for better non-wage job characteristics or amenities when making firm choices. Thus, firm-specific wage premiums could reflect unfavorable amenities, not imperfect competition.

The primary goal of our paper is to address these challenges and quantify the importance of imperfect competition in the U.S. labor market by estimating the size of rents earned by American firms and workers from ongoing employment relationships. To this end, we construct a matched employer-employee panel data set by combining the universe of U.S. business and worker tax records for the period 2001-2015. Using this panel data, we describe several important features of the U.S. labor market, including the size of firm-specific wage premiums, the sorting of workers to firms, the production complementarities between high ability workers and productive firms, and the pass through of firm and market shocks to workers' wages. Guided by these empirical results, we develop, identify and estimate a model of the labor market that allows us to draw inference about imperfect competition, compensating differentials and rent sharing. We also use the model to quantify the relevance of non-wage job characteristics and imperfect competition for inequality and tax policy, to assess the economic determinants of worker sorting, and to offer a unifying explanation of key empirical features of the U.S. labor market.

In Section 2, we describe the business and worker tax records, which provides us with panel data on the outcomes and characteristics of U.S. firms and workers. The firm data contain information on sales, revenues and inputs as well as industry codes and geographical identifiers. We merge the firm data set with worker tax records, creating the matched employer-employee panel data. In Section 3, we use this panel data data to describe key features of the U.S. labor

¹For a recent review of the literature, see Manning (2011).

market. This description is based on estimates from a statistical model of earnings that allows us to control for time-invariant unobserved heterogeneity of workers and firms while examining both the pass through of changes in firm performance to the earnings of its employees as well as the determinants of the cross-sectional distribution of earnings.

The pass through analysis follows the approach of Guiso et al. (2005), except we allow the earnings of an incumbent worker to respond differently to an idiosyncratic value added shock to the current firm than to a (same size) shock to all firms in the same local labor market. The analysis of cross-sectional earnings inequality builds and extends on the log additive worker and firm effects wage model proposed by Abowd et al. (1999, AKM hereafter). Unlike AKM, we allow for interactions between worker and firm effects and we let firm effects evolve over time as firm shocks pass through to workers' earnings. In contrast to much of the existing work, we also attempt to address the concern that estimates of firm effects will be biased upward and estimates of worker sorting will be biased downward in finite samples, with the size of the bias depending inversely on the degree of worker mobility among firms (Andrews et al., 2008).

The findings from the statistical model of earnings may be summarized with six broad conclusions. First, most of the variation in earnings is explained by heterogeneity in the quality of the workers as measured by their fixed effects. Second, firm effects explain only a few percent of the variation in earnings. Third, a substantial share of the variation in earnings is due to positive sorting of better workers to higher paying firms. Our preferred estimates suggest a correlation between worker effects and firm effects of about 0.4. Fourth, the gains in log earnings from moving to higher paying firms are considerably larger for better workers. These positive interaction effects are especially pronounced at the upper tails of the distributions of worker and firm effects. Fifth, idiosyncratic value added shocks to a firm transmit significantly to the earnings of its employees. Controlling for time-invariant firm and worker heterogeneity through a difference-in-differences strategy, we estimate that a 10 percent increase in the value added of a firm leads to a 1.4 percent increase in the earnings of incumbent workers. Lastly, the earnings of incumbent workers respond significantly less to an idiosyncratic value added shock to the current firm than to a (same size) shock to all firms in the same industry and location.

Motivated and guided by these findings, we develop in Section 4 an equilibrium model of the labor market that builds on work by Rosen (1986), Boal and Ransom (1997), Bhaskar et al. (2002), Manning (2003), and Card et al. (2018). Competitive labor market theory requires firms to be wage takers so that labor supply to the individual firm is perfectly elastic. The evidence that idiosyncratic productivity shocks to a firm transmit to the earnings of its workers is at odds with this theory. To allow labor supply to be imperfectly elastic, we let multiple employers compete with one another for workers who have heterogeneous preferences over amenities. Since we allow these amenities to be unobserved to the analyst, they can include a wide range of characteristics, such as distance of the firm from the worker's home, flexibility in the work schedules, the type of tasks performed, the effort required to perform these tasks, the social environment in the workplace, and so on.²

 $^{^{2}}$ There is limited empirical evidence on which non-wage characteristics matter the most. However, survey data from Maestas et al. (2018) point to the importance of flexibility in work schedules, the type of tasks performed,

The importance of workplace amenities has long been recognized in the theory of compensating differentials (Rosen, 1986). This is a theory of vertical differentiation: some employers offer better amenities than others. Employers that offer favorable amenities attract labor at lower than average wages, whereas employers offering unfavorable amenities pay premiums as offsetting compensation in order to attract labor. Our model combines this vertical differentiation with a horizontal employer differentiation: workers have different preferences over the same workplace amenities. As a result of this preference heterogeneity, the employer faces an upward sloping supply curve for labor, implying that wages are an increasing function of firm size. We assume that employers do not observe the idiosyncratic taste for amenities of any given worker. This information asymmetry implies that employers cannot price discriminate with respect to workers' reservation values. Instead, if a firm becomes more productive and thus wants to increase its size, the employer needs to offer higher wages to all workers of a given type. As a result, the equilibrium allocation of workers to firms creates surplus or rents to inframarginal workers.

The size of rents is determined by the slope of the labor supply curve facing the firm. The steeper the labor supply curve, the more important amenities are for workers' choices of firms as compared to wages. Therefore, imperfect competition as measured by rents increases in the progressivity of labor income taxes and in the variability of the idiosyncratic taste for amenities. However, the existence of rents does not imply the equilibrium allocation of workers is inefficient. In our model, the market allocation will be inefficient if the firms differ in wage setting power, and, thus, in the ability to mark down wages relative to the marginal product of labor. To allow for such differences, we let workers view firms as closer substitutes in some markets than others. This structure on the workers' preferences captures that workplace characteristics are likely to vary systematically across firms depending on location and industry.

In Section 5, we take the model to the data. To increase our confidence in the empirical findings from the model, we allow for rich unobserved heterogeneity across workers with respect to preferences and productivity and between firms in terms of technology and amenities. Even so, it is possible to prove identification of the parameters of interest given the panel data of workers and firms. To this end, we first show the restrictions on the model primitives that deliver the statistical model of earnings used to describe the labor market. This forges a direct link between our model of the labor market in Section 4 and the panel data analyses in Section 3. Once this link is established, we proceed by showing how the estimates produced by the statistical model allow us to recover the model-based quantities of interest. For example, the rents earned by workers can be measured by the elasticity of the labor supply curve to the firm. This elasticity can be recovered from the difference-in-differences estimate of the pass through of firm shocks to incumbent workers' earnings. Another example is the correlation structure in a worker's taste for amenities, which can be identified by comparing difference-in differences estimates of the pass-through rates of firm versus market level shocks. The estimates of worker effects, firm effects and worker sorting are also important for identification, allowing us to recover the productivity

and the amount of effort required. The analysis of Marinescu and Rathelot (2018) suggests distance of the firm from the workers' home may be important.

of workers, the compensating differentials due to the vertical differentiation of firms, and the extent to which preferences for amenities vary by worker productivity. To determine whether productive workers and firms are complements, we take advantage of the estimated interaction coefficients between worker and firm effects.

The model yields four key findings that we discuss in Section 6. First, there is a significant amount of rents and imperfect competition in the U.S. labor market due to horizontal employer differentiation. Workers are, on average, willing to pay 14 percent of their wage to stay in the current jobs. Comparing these worker rents to those earned by employers suggests that total rents are divided relatively equally between firms and workers. Second, the evidence of small firm effects do not imply that labor markets are competitive or that rents are negligible. Instead, firm effects are small because productive firms tend to have good amenities, which pushes down the wages that these firms have to pay. As a result of these compensating differentials, firms contribute much less to earnings inequality than what is predicted by the variance of firm productivity only. Third, the primary reason why better workers are sorting into better firms is production complementarities, not heterogeneous tastes for workplace amenities. These complementarities are key to explain the significant misallocation of workers to firms. We estimate that a tax reform which would eliminate labor and tax wedges would increase total welfare by 5 percent and total output by 3 percent.

Our paper relates to a considerable, but largely theoretical, literature on imperfect competition and rents in the labor market, reviewed in Boal and Ransom (1997) and Manning (2011). While our model creates wage setting power and rents from horizontal employer differentiation, alternative micro-foundations of the same imperfectly elastic labor supply curve may be possible. In such cases, the conclusions drawn about imperfect competition and rent sharing do not necessarily change. As noted by Manning (2011), the type of monopsony model that we use can generate a positive relationship between firm-specific productivity and wages that mimics models in which workers bargain with employers over wages. But in our model, firms set wages unilaterally and rents are shared only because of information asymmetries. Distinguishing between alternative sources of imperfect competition and rent sharing is interesting but beyond the scope of our paper. However, Card et al. (2018) argue that a monopsony type of explanation for rent sharing is more plausible than one based on worker bargaining power, particularly in countries that lack strong unions. One interesting implication of the monopsony explanation for rents is that there is no hold-up problem in the firm's investment decision. Card et al. (2013a) find support for this prediction in the data.

The insights from our paper contribute to a large and growing literature on firms and labor market inequality, reviewed in Card et al. (2018). A number of studies show that trends in wage dispersion closely track trends in productivity dispersion across industries and workplaces (Faggio et al., 2010; Dunne et al., 2004; Barth et al., 2016). While this correlation might reflect that some of the productivity differences across firms spill over to wages, it could also be driven by changes in the degree to which workers of different quality sort into different firms (see e.g. Murphy and Topel, 1990; Gibbons and Katz, 1992; Abowd et al., 1999; Gibbons et al., 2005). To address the sorting issue, a growing body of work has taken advantage of matched employer-employee data. Some studies use this data to estimate the pass through of changes in the value added of a firm to the wages of its workers, while controlling for time-invariant firm and worker heterogeneity.³ These studies typically report estimates of pass-through rates in the range of 0.05-0.20. We complement this work by providing evidence of the pass-through rates for a broad set of firms in the U.S. and by showing how the estimated pass through of firm and market level shocks can be used to draw inferences about imperfect competition, rents, and allocative inefficiency.

Another set of studies use the matched employer-employee data to estimate the additive worker and firm effects wage model proposed by AKM. These studies tend to conclude that firms play an important role in the wage determination, with a typical finding that about 15-20 percent of the variance of wages is attributable to firm effects (Card et al., 2018). Some studies interpret the sizable firm effects as evidence of rents and imperfect competition in the labor market. Our paper contributes to this literature in three ways. First, our study makes clear that firm effects neither imply nor are implied by rents. The relevant quantity to measure imperfect competition and rents is the sensitivity of wages to changes in firm productivity, not the firm effects. Second, we show that firm effects are small in the U.S. labor market, explaining only a few percent of the variation in earnings. This finding contrasts with recent work from the U.S. (Sorkin, 2018; Song et al., 2018) as well as many studies from other developed countries (Card et al., 2018). The reason is that these studies do not address the concern that estimates of firm effects will be biased upward and estimates of worker sorting will be biased downward due to limited worker mobility across firms (Andrews et al., 2008). Following the approaches of Bonhomme et al. (2019, BLM hereafter) and Kline et al. (2018b), we try to correct for these biases. Both approaches show that firm effects explain very little of the variation in earnings in the U.S. economy, once one corrects for bias due to limited mobility. Instead, a substantial part of the variation in earnings is due to positive sorting of high wage workers to high paying firms. These findings conform to what BLM find for Sweden. Third, we find evidence of non-additivity in the U.S. wage structure. The interaction effects are economically significant, incentivizing better workers to sort into productive firms.

Our paper also relates to a literature that tries to measure the role of compensating differentials for wage setting and earnings inequality. This literature is reviewed in Taber and Vejlin (2016) and Sorkin (2018). Much of the existing evidence comes from hedonic regressions of earnings on one or more observable non-wage characteristics of jobs, employers, or industries, interpreting the regression coefficients as the market prices of those amenities. Typical estimates of these coefficients are small in magnitude and sometimes of the wrong sign (Bonhomme and

³See e.g. Guiso et al. (2005), Card et al. (2013a), Card et al., 2018, Carlsson et al. (2016), Lamadon (2016), Friedrich et al. (2019). A concern with this approach is that measures of firm productivity may reflect a number of factors. Some studies have therefore examine the pass through of specific, observable changes. For example, Van Reenen (1996) studies how innovation affects firms' profit and workers' wages. He also investigated patents as a source of variation, but found them to be weakly correlated with profits. Building on this insight, Kline et al. (2018a) studies the incidence of patents that are predicted to be valuable. A related literature has examined the wage and productivity effects of adoption of new technology in firms (see Akerman et al., 2015, and the references therein).

Jolivet, 2009). However, these estimates could be severely biased, either due to correlations between observed amenities and unobserved firm characteristics or because of assortative matching (on unobservables) between workers and firms (see e.g. the discussion in Ekeland et al., 2004). Several recent studies have used panel data in an attempt to address these concerns. Like us, Taber and Vejlin (2016), Lavetti and Schmutte (2017), and Sorkin (2018) take advantage of matched longitudinal employer-employee data to allow for unobserved heterogeneity across firms.

Our paper differs from the existing literature on compensating differentials in several ways. One important difference is that amenities, in our model, create both vertical and horizontal employer differentiation. The latter generate imperfect competition, wage setting power and rents; the former acts as standard compensating differentials. By comparison, compensating differentials have typically been analyzed in models with perfect competition or search frictions (see e.g. Mortensen, 2003). Our paper also allows for ex-ante worker heterogeneity in productivity and preferences which generate sorting between firms and workers, in contrast to, for example, Sorkin (2018). Our estimates suggest that worker heterogeneity and sorting are empirically important features of the U.S. labor market which are necessary to take into account to understand the determinants of earnings inequality. By taking our model to the data, we are able to quantify the relative importance of amenities versus production complementarities for worker sorting and earnings inequality. Lastly, our paper differs in that we move beyond the impact of amenities on wages and worker sorting, examining also the implications for tax policy and allocative efficiency. In our model, wages are taxed but the (idiosyncratic taste for) amenities are not. Thus, progressive taxation on labor income may distort the worker's decision of which firm and market to work in. We analyze, theoretically and empirically, the consequences of this distortion and how changes in the tax system may help improve the allocation of workers to firms.⁴

2 Data Sources and Sample Selection

2.1 Data sources

Our empirical analyses are based on a matched employer-employee panel data set with information on the characteristics and outcomes of U.S. workers and firms. This data is constructed by linking U.S. Treasury business tax filings with worker-level filings for the years 2001-2015. Below, we briefly describe data sources, sample selection, and key variables, while details about the data construction and the definition of each of the variables are given in Online Appendix A.

Business tax returns include balance sheet and other information from Forms 1120 (C-corporations), 1120S (S-corporations), and 1065 (partnerships). The key variables that we

 $^{^{4}}$ Tax theory in the Mirrlees (1971) tradition generally assumes the labor markets are perfectly competitive. A notable exception is Cahuc and Laroque (2014) who develop a model for optimal taxation under monopsonistic markets. A larger literature considers tax design in situations with search frictions. See Yazici and Sleet (2017) and the references therein.

draw on from the business tax filings are the firm's employer identification number (EIN) and its value added, commuting zone, and industry code. Value added is the difference between receipts and the cost of goods sold. Commuting zone is constructed using the ZIP code of the firm's business filing address. Industry is defined as the first two digits of the firm's NAICS code. We define a market as the combination of an industry and a commuting zone. At times we will aggregate these markets according to the combination of Census regions (Midwest, Northeast, South, West) and broad sectors (Goods and Services). We will refer to this classification as broad markets.

Earnings data are based on taxable remuneration for labor services reported on form W-2 for direct employees and on form 1099 for independent contractors. Earnings include wages and salaries, bonuses, tips, exercised stock options, and other sources of income deemed taxable by the IRS. These forms are filed by the firm on behalf of the worker and provide the firmworker link. We express all monetary variables in 2015 dollars, adjusting for inflation using the Consumer Price Index.

2.2 Sample Selection

In each year, we start with all individuals aged 25-60 who are linked to at least one employer. Next, we define the worker's firm as the EIN that pays her the greatest direct (W-2) earnings in that year. This definition of a firm conforms to previous research using the U.S. business tax records (see, e.g., Song et al., 2018). The EIN defines a corporate unit for tax and accounting purposes. It is a more aggregated concept than an establishment, which is the level of analysis considered in recent research on U.S. Census data (see, e.g., Barth et al., 2016), but a less aggregated concept than a parent corporation. As a robustness check, we investigated the sensitivity of the estimated firm wage premiums to restricting the sample to EINs that appear to have a single primary establishment. These are EINs for which the majority of workers live in the same commuting zone. It is reassuring to find that the estimated firm wage premiums do not materially change when we use this restricted sample.⁵

Since we do not observe hours worked or a direct measure of full-time employment, we follow the literature by including only workers for whom annual earnings are above a minimum threshold (see, e.g., Song et al., 2018). In the baseline specification, this threshold is equal to \$15,000 per year (in 2015 dollars), which is approximately what people would earn if they work full-time at the federal minimum wage. As a robustness check, we investigate the sensitivity of our results to other choices of a minimum earnings threshold. We further restrict the sample to firms with non-missing value added, commuting zone, and industry. The full sample includes 447.5 (39.2) million annual observations on 89.6 (6.5) million unique workers (firms).

In parts of the analysis, we consider two distinct subsamples. The first subsample, which we refer to as the *stayers sample*, restricts the full sample to workers observed with the same employer for eight consecutive years. This restriction is needed to allow for a flexible specification

 $^{^{5}}$ In the baseline sample, the AKM (BLM) estimates of firm effects are around 10 (3) percent. By comparison, the restricted sample gives AKM (BLM) estimates of approximately 9 (3) percent.

		Workers		Firms	
Panel A.		Baseline	Sample		
	Unique	Observation-Years	Unique	Observation-Years	
Full Sample:	$89,\!570,\!480$	$447,\!519,\!609$	$6,\!478,\!231$	$39,\!163,\!975$	
Panel B.		Movers	Sample		
	Unique	Observation-Years	Unique	Observation-Years	
Movers Only:	$32,\!070,\!390$	$207,\!990,\!422$	$3,\!559,\!678$	$23,\!321,\!807$	
Panel C.		Stayers Sample			
	Unique	6 Year Spells	Unique	6 Year Spells	
Complete Stayer Spells:	$10,\!311,\!339$	$35,\!123,\!330$	$1,\!549,\!190$	$6,\!533,\!912$	
10 Stayers per Firm:	$6,\!297,\!042$	$20,\!354,\!024$	144,412	$597,\!912$	
10 Firms per Market:	$5,\!217,\!960$	$16,\!506,\!865$	$117,\!698$	476,878	

Table 1: Overview of the Sample

Notes: This table provides an overview of the full sample, movers sample, and stayers sample, including the steps involved in defining the stayers sample.

of how the worker's earnings evolve over time. Specifically, we omit the first and last years of these spells (to avoid concerns over workers exiting and entering employment during the year, confounding the measure of annual earnings) and analyze the remaining six-year spells. Furthermore, the stayers sample is restricted to employers that do not change commuting zone or industry during those eight years. Lastly, we restrict the stayers sample to firms with at least 10 such stayers and markets with at least 10 such firms, which helps to ensure sufficient sample size to perform the analyses at both the firm and the market level. The stayers sample includes 35.1 (6.5) million spells on 10.3 (1.5) million unique workers (firms).

The second subsample, which we refer to as the *movers sample*, restricts the full sample to workers observed at multiple firms. That is, it is not the same EIN that pays the worker the greatest direct (W-2) earnings in all years. Following previous work, we also restrict the movers sample to firms with at least two movers. This restriction might help reduce the limited mobility bias. It also makes it easier to directly compare the AKM and BLM estimates of firm effects to those produced by the approach of Kline et al. (2018b) (which requires at least two movers per firm). The movers sample includes 32.1 (3.6) million unique workers (firms).

Table 1 compares the size of the baseline, the stayers, and the movers samples. Detailed summary statistics of these samples of linked firms and worker are given in Online Appendix Table A.1. The samples are broadly similar, both in the distribution of earnings but also in firm-level variables such as value added, wage bill, size, and the geographic distribution across regions and sectors. The most noticeable differences are that the stayers have, on average, somewhat higher earnings and tend to work in firms with higher value added.

3 Key features of the U.S. labor market

In this section, we use the panel data on workers and firms to describe key features of the U.S. labor market. We begin by describing the statistical model of earnings that we will apply to this

data. Next, we present the empirical findings, and then discuss how they motivate and guide our choices of how to model the labor market.

3.1 Statistical model of earnings

We assume that workers' earnings can be described by the following equation:

$$\log W_{it} = \mathcal{X}'_{it}\vartheta + w_{it},\tag{1}$$

where W_{it} denotes the earnings for individual *i* in year *t*, \mathcal{X}_{it} is a vector of covariates which includes a full set of indicators for calendar years and a cubic polynomial in age, and w_{it} denotes log earnings net of age effects and common aggregate time trends. As described below, we allow w_{it} to depend on both the workers' own productivity and the firm in which she works. Our measure of firm performance is value added, which is determined by the equation:

$$\log Y_{jt} = \mathcal{Z}'_t \varphi + y_{jt},\tag{2}$$

where Y_{jt} denotes the value added for firm j in year t, Z_t includes a full set of indicators for calendar years, and y_{jt} is log value added net of common aggregate time trends. The key elements of equations (1) and (2) are the time series properties of w_{it} and y_{jt} , which we now specify.

3.1.1 Specification of processes

We assume that y_{jt} evolve according to the following process:

$$y_{jt} = \zeta_j + y_{jt}^p + \xi_{jt} + \delta^y \xi_{jt-1}$$
(3)

$$y_{jt}^p = y_{jt-1}^p + u_{jt},$$
(3)

$$u_{jt} = \tilde{u}_{jt} + \bar{u}_{r(j),t}$$

where r(j) denotes the market of firm j, ζ_j is a fixed effect for the firm, and the time-varying part of y_{jt} is decomposed into a permanent component, assumed to follow a unit root process with innovation shock u_{jt} , and a transitory component, which is assumed to follow a MA(1) process with coefficients δ^y and innovation variance σ_{ξ}^2 . The permanent innovation u_{jt} consists of a common innovation to all firms in a given market r, $\bar{u}_{r(j),t} \equiv \mathbb{E}[u_{jt}|r(j)=r]$, and an idiosyncratic innovation specific to the firm, $\tilde{u}_{jt} \equiv u_{jt} - \bar{u}_{r(j),t}$.

We assume that w_{it} evolve according to the following process:

$$w_{it} = \phi_{ij(i,t)} + w_{it}^{p} + \nu_{it} + \delta^{w} \nu_{it-1}$$

$$w_{it}^{p} = w_{it-1}^{p} + \gamma \tilde{u}_{j(i,t),t} + \Upsilon \bar{u}_{r(i,t),t} + \mu_{it},$$
(4)

where j(i, t) and r(i, t) denote the firm and market of worker *i* in year *t*, and ϕ_{ij} is a fixed effect for worker *i* if she works in firm *j*. The time-varying part of w_{it} is decomposed into a permanent component w_{it}^p and a transitory component, assumed to follow a MA(1) process with coefficients δ^w and innovation variance σ_v^2 . The permanent earnings component evolves for three reasons: worker-specific innovations μ_{it} , pass through of firm-specific value added shocks $\gamma \tilde{u}_{j(i),t,t}$, and pass through of market level value added shocks $\gamma \bar{u}_{r(i,t),t}$.

3.1.2 Parameters of interest and assumptions

Our interest is centered on two aspects of this statistical model of earnings. The first is how changes in firm performance affect the earnings of incumbent workers, as measured by the passthrough rates γ and Υ . The second is the determinants of the cross-sectional distribution of earnings, which we measure by decomposing ϕ_{ij} into components that capture worker heterogeneity, firm-specific wage premiums, worker sorting, and interactions between worker and firm effects.

For these purposes, it is necessary to invoke some restrictions on the statistical model of earnings. Let $J = \{j(i,t)\}_{i,t}$ and $U = \{\tilde{u}_{jt}, \bar{u}_{r(j),t}\}_{j,t}$ and $Q = \{\xi_{jt}\}_{j,t}$. We make the following assumptions:

Assumption 1. $\mathbb{E}[\xi_{jt}|r(j)=r, J, U] = \mathbb{E}[\xi_{jt'}\xi_{jt}|r(j)=r, J, U] = 0$ for all j, r, t, t'.

Assumption 2. $\mathbb{E}[\mu_{it}, \nu_{it}|J, U, Q] = 0$ for all i, t.

Assumption 1 is the same restriction on the error structure of the value added process as in Guiso et al. (2005). It implies that transitory shocks to value added are mean zero and uncorrelated with past transitory shocks to value added. Assumption 2 is a condition on the relationship between the worker-specific innovations to earnings, worker mobility, and innovations to firm value added. The assumption embodies two types of economic restrictions. The first restriction, from conditioning on j(i,t), implies that mobility is exogenous to the worker-specific innovations to earnings (which are paid to the worker independent of the choice of firm). This is the same restriction on worker mobility as invoked in the Abowd et al. (1999) model. The second restriction, from the conditioning on the innovations to firm value added, implies that the worker-specific innovations to earnings neither co-vary across coworkers nor with shocks to firm value added. This is the same restriction as in Guiso et al. (2005).

It is important to observe what is *not* being restricted under Assumptions 1 and 2. First, we do not restrict whether or how workers sort into firms according to the worker effects, the firm effects, or the interactions between the worker and firm effects. Second, we do not restrict whether or what type of workers move across firms in response to innovations to firm value added. In fact, workers with different values of ϕ_{ij} may have arbitrarily different mobility patterns. Third, the statistical model of earnings does not specify why individuals choose the firm that they do. However, it also does not preclude the possibility that individuals choose firms to maximize earnings or utilities. For instance, Assumptions 1 and 2 are consistent with each worker choosing the firm that offers his preferred combination of wages and non-wage attributes, as shown in Section 5.

3.2 Pass through of firm shocks

In this section, we are interested in estimating the parameters γ and Υ , which we refer to as the *pass-through rates* of firm-specific and market level value added shocks. Before presenting estimates of the pass-through rates, we show how these parameters can be identified through a difference-in-differences (DiD) strategy.

3.2.1 Identification, moment conditions and DiD representation

To compare with existing work, we first consider a special case of the statistical model of earnings where $\gamma = \Upsilon$. That is, we assume the pass-through rate of an idiosyncratic value added shock to the current firm is of the same size as the pass-through rate of a value added shock to all firms in the current market. We focus on the sample of stayers as captured by the indicator variable $S_i = 1[j(i, 1) = ... = j(i, T)].$

As shown in Lemma 1 in Online Appendix B.5, Assumptions 1 and 2 give the following moment conditions:

$$\mathbb{E}\left[\Delta y_{j(i)t}\left(w_{it+\tau} - w_{it-\tau'} - \gamma\left(y_{j(i),t+\tau} - y_{j(i),t-\tau'}\right)\right) | S_i = 1\right] = 0$$
for $\tau \ge 2, \tau' \ge 3$

$$(5)$$

Solving for γ we identify the pass through of a firm-specific shock to the earnings of incumbent workers:

$$\gamma = \frac{\mathbb{E}\left[\Delta y_{j(i)t} \left(w_{it+\tau} - w_{it-\tau'}\right) | S_i = 1\right]}{\mathbb{E}\left[\Delta y_{j(i)t} \left(y_{j(i),t+\tau} - y_{j(i),t-\tau'}\right) | S_i = 1\right]}$$

Thus, we can identify the pass through of a firm-specific shock from our panel data on firms and workers.

DiD interpretation

To interpret this identification result and assess the underlying assumptions, note that the statistical model of earnings includes fixed effects for time and agents. By controlling for these fixed effects we obtain a DiD strategy, looking within workers and firms while eliminating common changes over time in the labor market or the economy more generally. To see the DiD representation, suppose for simplicity the workers can be assigned to two groups of firms: one half has $\Delta y_{j(i)t} = +\delta$ and the other half has $\Delta y_{j(i)t} = -\delta$. We then get the following interpretation of γ as the ratio of two DiDs.

$$\gamma = \frac{\mathbb{E}\left[w_{it+\tau} - w_{it-\tau'} | +\delta, S_i=1\right] - \mathbb{E}\left[w_{it+\tau} - w_{it-\tau'} | -\delta, S_i=1\right]}{\mathbb{E}\left[y_{j(i),t+\tau} - y_{j(i),t-\tau'} | +\delta, S_i=1\right] - \mathbb{E}\left[y_{j(i),t+\tau} - y_{j(i),t-\tau'} | -\delta, S_i=1\right]}$$

Under an assumption of common underlying trends between the two groups, the numerator gives the treatment effect on log earnings; the denominator gives the treatment effect on log value added; and the ratio gives the elasticity of earnings with respect to value added.

Graphical evidence

In Figure 1, we empirically assess the DiD strategy. The figure is constructed in the following way: In any given calendar year (denoted period t = 0), we i) order firms according to the increase $\Delta y_{j(i)t}$; ii) separate the firms at the median in the distribution of $\Delta y_{j(i)t}$, letting the upper half constitute the treatment firms and the lower half the control firms; and iii) plot the differences in y_{jt} between these two groups in period t = 0 as well as in the years before (periods t < 0) and after (periods t > 0). We perform these three steps separately for various calendar years, always weighting each firm by the number of workers. The solid (dashed) black line represents the difference in log value added (wages) for the treatment and control firms where each firm is weighted by the number of workers.

By construction, the treatment and control groups differ in the value added growth from period t-1 to period t. On average, firms in the treatment group experience about 30 percentage points larger growth in value added as compared to firms in the control group. According to the value added process (3), the growth in value added should be the sum of a permanent component and a transitory, mean-reverting component. Due to the transitory component, $\Delta y_{j(i)t}$ could be correlated with $\Delta y_{j(i)\tau}$ at $\tau = t - 2, ..., t + 2$. However, $\Delta y_{j(i)t}$ should be orthogonal to $\Delta y_{j(i)\tau}$ in the periods before $\tau = t - 2$ and after $\tau = t + 2$. Consistent with this orthogonality condition, the figure shows a very similar trend in log value added between the treatment and control group at these periods. Reassuringly, firms that experienced large growth in value added in period 0 are no more or less likely to experience large growth in value added in periods -6 to -3 or in periods 3 to 6.

The dashed black line performs the same exercise, but this time for log wages of incumbent workers who stay in the firm in all six years. On average, workers in treatment firms experience an additional 5 percentage points increase in earnings in period 0 as compared to workers in the control firms. Interpreted through the lens of the DiD design, this finding suggests a passthrough rate of firm shocks γ above .15. The growth in earnings is also the sum of a permanent component and a transitory, mean-reverting component. Therefore, Δw_{it} could be correlated with $\Delta w_{i\tau}$ at $\tau = t - 2, ..., t + 2$, but it should be orthogonal to $\Delta w_{i\tau}$ in the periods before $\tau = t - 2$ and after $\tau = t + 2$. Reassuringly, the dashed line shows a very similar trend in log earnings between workers in the treatment and control group during these periods.

Firm versus market level shocks

We now shift attention to the general case where γ may differ from Υ , thereby allowing the earnings of an incumbent worker to respond differently to an idiosyncratic value added shock to the current firm than to a (same size) shock to all firms in a given market. To identify the firm-level pass-through rate γ , we then need to demean the variables of interest using a within market times year transformation, $\tilde{w}_{it} = w_{it} - \mathbb{E}[w_{it}|r(i,t)=r]$ and $\tilde{y}_{jt} = y_{jt} - \mathbb{E}[y_{jt}|r(j)=r]$. As shown in Online Appendix B.5, Assumptions 1 and 2 then give the following moment conditions that we can use to identify the pass-through rates of firm-specific value and market level added



Figure 1: Difference-in-differences representation of the estimation procedure

Notes: This figure displays the mean differences in log value added (solid lines) and log earnings (dotted lines) between firms that receive an above-median versus below-median log value added change at event time zero. Results are presented for the unconditional measures of log value added and log earnings (black lines), for the measures of log value added and log earnings net of market interacted with year effects (red lines), and for the averages of log value added and log earnings by market and year (blue lines). The shaded area denotes the time periods during which the orthogonality condition need not hold in the identification of the permanent pass-through rate.

shocks by solving for γ and Υ :

$$\mathbb{E}\left[\Delta \tilde{y}_{j(i),t}\left(\tilde{w}_{it+\tau} - \tilde{w}_{it-\tau'} - \gamma\left(\tilde{y}_{j(i),t+\tau} - \tilde{y}_{j(i),t-\tau'}\right)\right)|S_i=1\right] = 0$$
(6)

$$\mathbb{E}\left[\Delta \bar{y}_{j(i),t}\left(\bar{w}_{it+\tau} - \bar{w}_{it-\tau'} - \Upsilon\left(\bar{y}_{j(i),t+\tau} - \bar{y}_{j(i),t-\tau'}\right)\right)|S_i=1\right] = 0$$
(7)

for
$$\tau \ge 2, \tau' \ge 3$$

where $\bar{y}_{r(j),t} \equiv \mathbb{E}[y_{jt}|r(j)=r]$ and $\bar{w}_{r(j),t} \equiv \mathbb{E}[w_{jt}|r(j)=r]$.

The red and blue lines in Figure 1 represent the differences between the treatment and control group in $(\tilde{w}_{it}, \tilde{y}_{jt})$ and $(\bar{w}_{r(i,t),t}, \bar{y}_{r(j),t})$ over time. These lines are constructed in the same way as the black lines, except the red and blue lines use the demeaned variables \tilde{w} and \tilde{y} and the market averages \bar{w} and \bar{y} , respectively. Comparing the red solid line to the red dashed line reveals that conditioning on the full set of year times market fixed effects attenuates slightly the treatment effect on log earnings relative to the treatment effect on log value added. Interpreted through the lens of the DiD design, this finding suggests the estimated pass-through rate of a firm-specific shock will be slightly lower once we allow for Υ to differ from γ . By way of comparison, the DiD applied to the market averages of wages and value added suggests a relatively large estimate of Υ . Thus, we expect the estimated pass-through rate of an idiosyncratic value added shock to the current firm γ to be smaller than the pass-through rate of a same size shock to all firms in

the market Υ .

3.2.2 Estimates of the pass-through rates

To obtain point estimates and standard errors of the pass-through parameters, we jointly estimate the earnings and value added processes (3) and (4) using generalized method of moments (GMM). These processes are estimated on the stayers sample, pooling the data for all years but taking out common calendar year effects. The parameter estimates are summarized in Table 2. We refer to Online Appendix B.5 for the full set of parameter estimates and for details on how the processes are estimated. Specification checks are reported in Online Appendix Table B.1, supporting our main specification of the time series properties of w_{it} and y_{jt} .

Main results

To compare with existing work, we first consider the estimation results from the restricted specification that imposes $\Upsilon = \gamma$. These results are reported in the first two columns of Table 2. The standard deviation in log earnings growth is 0.17. Decomposing the variance in log earnings growth, we find that almost 40 percent is due to permanent shocks at the worker level, 58 percent can be attributed to transitory shocks at the worker level, and 3 percent is due to the pass through of permanent shocks to value added at the firm level. The estimated pass-through rate $\hat{\gamma}$ is 0.14 with a standard error of 0.01, suggesting that a 10 percent permanent increase in the value added of the firm leads to a 1.4 percent permanent increase in the earnings of incumbent workers.

As shown in the last two columns of Table 2, however, the estimated pass through differs materially depending on whether the shock is specific to the firm or common to all firms in the same market. Conditional on the full set of year times market fixed effects, the pass-through rate of a firm-level shock is very precisely estimated at 0.13. By way of comparison, a common shock to firms in the same market has a much larger pass through. The estimated market passthrough rate \hat{T} is 0.18 with a standard error of 0.02, suggesting that the earnings of incumbent workers increases by 1.8 percent if all firms in their market experience a 10 percent permanent increase in value added. This finding highlights the importance of distinguishing between shocks that are specific to workers in a given firm versus those that common to workers in a market.

Our specification of the earnings process allows permanent shocks to value added to be transmitted to workers' earnings, whereas transitory firm shocks are not. As a specification check, we allow transitory innovations to value added to transmit to workers' earnings. We find little if any pass through of transitory shocks. As a result, transitory shocks explain as little as 0.1 percent of the variation in log earnings. This finding is consistent with previous work (see e.g. Guiso et al. 2005; Friedrich et al. 2019). A possible interpretation of this finding is that transitory changes in value added reflect measurement error that do not give rise to economic responses. In the remainder of the paper, we will treat the transitory changes in value added as measurement error and focus on the pass through of the permanent shocks.

	Paramete	ers and Gr	owth Decon	nposition
	Firm	Only	Accounting	for Markets
	Parameter	Var. (%)	Parameter	Var. (%)
Permanent Worker Shock (Std. Dev.)	0.10 (0.00)	39.5%	0.10 (0.00)	38.1%
Transitory Worker Shock (Std. Dev.)	0.13 (0.00)	57.6%	$\begin{array}{c} 0.13 \\ (0.00) \end{array}$	57.4%
Permanent Firm Shock Passed-through (Std. Dev.)	0.03 (0.00)	2.8%	0.02 (0.00)	1.8%
— Permanent Firm Shock Passthrough Coefficient	0.14 (0.01)		0.13 (0.01)	
Transitory Firm Shock Passed-through (Std. Dev.)	0.00 (0.00)	0.0%	0.00 (0.00)	0.0%
— Transitory Firm Shock Passthrough Coefficient	-0.01 (0.01)		0.00 (0.00)	
Market Shock Passed-through (Std. Dev.)			0.02 (0.00)	1.1%
— Market Shock Passthrough Coefficient			0.18 (0.02)	

Table 2: GMM estimates of the earnings and value added processes

Notes: This table displays the pass-through rates and the decomposition of earnings growth of the joint processes of log value added and log earnings. These results come from joint estimation of the earnings and value added processes (3) and (4) using GMM. Columns 1-2 report results from the specification which imposes $\Upsilon = \gamma$ ("Firm only"), while columns 3-4 report results from the specification which allows Υ to differ from γ ("Accounting for Markets"). "Var. (%)" refers to the percent of log earnings changes explained. The variance explained by transitory worker shock refers to the share of the total variance explained by the MA(1) process. The variance explained by permanent firm (market) shock passed-through refers to the share of the total variance explained by permanent value added shocks at the firm (market) level that are passed through to workers' earnings. Standard errors are estimated using 40 block bootstrap draws in which the block is taken to be the market. See Online Appendix Table B.1 for the full set of parameter estimates (Panel A) and for results from the MA(2) specification of the transitory process (Panel B)

Heterogeneity and robustness

In Online Appendix Figure B.1, we explore heterogeneity and robustness of the pass through estimates. This figure reports separate estimates of the earnings and value added processes (3) and (4) for each subgroup of stayers. Conditional on a full set of year times market fixed effects, we find that the pass-through rates do not vary that much by the worker's age, previous wage, or gender. Moreover, the pass through rates do not change materially if we restrict the sample to new workers who were first hired at the firm in the beginning of the eight year employment spell versus those that have stayed in the firm for a longer time.

In Online Appendix Figure B.1b, we present results from several specification checks. Following Guiso et al. (2005), our main measure of firm performance is value added. They offer two reasons for using value added as a measure of firm performance. First, they argue, value added is the variable that is directly subject to stochastic fluctuations. Second, firms have discretionary power over the reporting of profits in balance sheets, which makes profits a less reliable objective to assess. Nevertheless, it is reassuring to find that the estimates of the pass-through rates are broadly similar if we measure firm performance by operating profits, earnings before interest, tax and depreciation (EBITD), or value added net of reported depreciation of capital. We also show that the estimated pass-through rate is in the same range as our baseline result if we exclude multinational corporations (for which it can be difficult to accurately measure value added) or exclude the largest firms (that are more likely to have multiple plants, which may not necessarily have the same wage setting).

3.3 Worker heterogeneity, firm wage premiums and worker sorting

We now turn attention to quantifying the importance of worker heterogeneity, firm wage premiums, worker sorting, and interactions between worker and firm effects as sources of inequality in the U.S. labor market.

3.3.1 Two-way (worker and firm) fixed effect model

To begin with, we consider a special case which assumes that $\phi_{ij} = x_i + \psi_j$ and that $\gamma = \Upsilon = 0$. The first restriction imposes a log additive structure on the earnings that worker *i* can expect to receive from working in firm *j*. Under this functional form, the worker fixed effect captures the (time-invariant) portable component of earnings ability, whereas the firm fixed effect can be interpreted as a firm-specific relative pay premium. The second restriction assumes there is no pass through of firm or market level shocks. As a result, the firm effects on earnings do not vary over time. By invoking these two restrictions, our statistical model of earnings reduces to the two-way (worker and firm) fixed effect model of AKM.

Under the above restrictions, the variance of log earnings can be written as:

$$Var(\log W_{it}) = \underbrace{Var(x_i + \mathcal{X}'_{it}b)}_{\text{Worker component}} + \underbrace{Var(\psi_{j(i,t)})}_{\text{Firm component}} + \underbrace{2Cov(x_i + \mathcal{X}'_{it}b, \psi_{j(i,t)})}_{\text{Sorting component}} + \underbrace{Var(\epsilon_{it})}_{\text{Residual}}$$
(8)

where the worker and firm components tell us how much of the variation in log earnings can be attributed to heterogeneity in worker and firm effects, respectively. The third component captures the contribution to earnings inequality from the sorting of workers to firms. The goal is to quantify these three components to draw inference about the determinants of earnings inequality in the U.S. economy. The decomposition includes both workers who move between firms and stayers. However, the firm and worker effects are only separately identified within a connected set of firms that are linked by worker mobility. Consistent with previous work, we therefore restrict our sample of workers (including stayers and movers) to those who work at a firm in the largest connected set in each time interval (2001-2008 and 2008-2015). In the U.S., this set covers more than 90 percent of the workers (see Online Appendix Table B.2).

Limited mobility bias

Even if the above restrictions hold, it is challenging to draw inference about the inequality contribution from firm effects and worker sorting. A key challenge is the incidental parameter bias caused by the large number of firm-specific parameters that are solely identified from workers who move across firms. The analysis of Andrews et al. (2008) suggests this limited mobility bias



Figure 2: Empirical Characterization of Limited Mobility Bias

Notes: In this figure, we consider the subset of firms with at least 15 movers. We randomly remove movers within each firm and re-estimate the variance of firm effects using the AKM and BLM estimators. For each estimator, we repeat this procedure several times, and then take averages of the variance estimates across these repetitions. The procedure allows us to keep the connected set of firms approximately the same and examine the bias that results from having fewer movers available in estimation.

can be substantial. With few movers per firm, the firm component is biased upwards while the sorting component is biased downwards, with the size of the bias depending inversely on the degree of worker mobility among firms.

To get a better sense of the scope for limited mobility bias in the U.S. data, we would ideally apply the AKM estimator to alternative samples of workers and firms that are comparable except for the number of movers per firm. Figure 2 presents the results from such an analysis, suggesting that the variance of firm effects declines monotonically as the number of movers per firm increases. To construct this figure, we consider a subsample of firms with reasonably many movers; that is, at least 15 movers per firm over the period 2001-2008. Applying AKM to this subsample gives an estimate of the variance of firm effects of 6.7 percent. Next, we remove movers randomly within firms (keeping the connected set of firms approximately the same) before re-estimating the AKM model. The solid line displays the AKM estimates of the variance of firm effects after randomly removing movers. Consistent with limited mobility bias, the fewer the number of movers per firm, the larger the variance of firm effects. For approximately the same set of firms, the estimated variance of firm effects is several times as large (23 percent) if we only keep ten percent of the movers within each firm (on average, 7 movers per firm) as compared to what we obtained if we keep all the movers per firm (at a minimum 15 and, on average, 62 movers per firm). By way of comparison, there are around 18 movers per firm in the full estimation sample (which roughly corresponds to the number of movers per firm when randomly removing 40% of movers).

Until recently, the procedures for addressing limited mobility bias required strong and questionable assumptions about the covariance structure of the time-varying errors (see e.g. the discussion in Card et al., 2018). To address this shortcoming, BLM and Kline et al. (2018b) propose approaches to address limited mobility bias that rely on a different or weaker set of assumptions.⁶ The first approach reduces the dimension of firm heterogeneity to a finite number of types. BLM show how this approach can be used to alleviate the biases arising from low mobility rates. The second approach uses a version of the Jackknife method. Kline et al. (2018b) show how this approach allows one to relax the homoskedasticity assumption in the bias correction procedure proposed by Andrews et al. (2008). Since it is computationally infeasible to apply Andrews et al. (2008) and Kline et al. (2018b) to very large data sets (as one needs to compute the trace of the inverse of the mobility matrix), our main analysis is based on the approach of BLM. As a robustness check, however, we use a subset of the U.S. states to assess the sensitivity of the results to the choice of procedure for addressing limited mobility bias.

In Figure 2, the dotted line shows estimates of the variance of firm effects based on the procedure of BLM that addresses limited mobility bias. As described in Section 5.4, firms are first classified into groups based on the empirical earnings distribution using the k-means clustering algorithm. The k-means classification groups together firms whose earnings distribution is most similar. Then, in a second step, the worker effects and firm effects are estimated. While the specification of BLM in Figure 2 assumes there exists 10 firm types, Online Appendix Figure B.4 shows the BLM estimates do not materially change if we instead allow for 20, 30, 40 or 50 firm types. Consistent with limited mobility bias, the BLM estimates are noticeably smaller than the standard AKM estimates in the samples with few movers. As expected, the AKM estimates become more similar to the BLM estimates when there is a large number of movers per firm, and thus, limited mobility bias should be small.

Estimates of worker effects, firm effects, and worker sorting

While the analysis in Figure 2 is useful to illustrate the scope for limited mobility bias, it does not offer estimates of firm effects for the entire connected set. In Table 3, we present results from the variance decomposition in (8) based on data for all firms and workers in the connected set (which includes both workers who move between firms and stayers). This table reports estimates of the worker, firm and sorting components as defined in equation (8).

Consider first Panel A of Table 3 where we present estimates from the AKM estimator for two different time periods (2001-2008 and 2008-2015) as well as pooled estimates where we combine the data from these time periods. The results show that the worker, firm and sorting components change little over time. Therefore, we focus attention on the pooled estimates. These results suggest that the firm effects explain around 9 percent of the variation in log earning, whereas worker sorting accounts for 5 percent. The correlation between firm effects and worker effects is only 0.1.

 $^{^{6}}$ Borovickova and Shimer (2017) propose a different approach that redefines firm types as the average firm wage rather than the wage premium, and relies on independence restrictions to recover the variance decomposition for this alternative definition.

Next, consider Panel B of Table 3 where we report the BLM estimates. As discussed above, a possible advantage of the BLM estimator is that it addresses limited mobility bias. Once we correct for such bias we find that firm effects are very small in the U.S. labor market, accounting for only 3 percent of the variation in log earnings. Instead, a larger part of the earnings variation is explained by worker sorting. The correlation between firm effects and worker effects exceeds 0.4 once we correct for limited mobility bias. This finding suggests that sorting of better workers to better firms is an empirically important feature of the U.S. labor market.

Our findings of small firm effects and strong sorting are at odds with recent work from the U.S. (Sorkin, 2018; Song et al., 2018) as well as many studies from other developed countries (Card et al., 2018). We argue the reason is that the existing literature do not properly address the concern over limited mobility bias.⁷ This raises questions such as: How do our results from the AKM estimator compare to those reported in the existing literature? Are the bias-corrected estimates sensitive to the procedure used?

To examine the first question of how our AKM results compare to existing work, consider Song et al. (2018, using SSA data from all U.S. states) and Sorkin (2018, using LEHD data for a subset of states). Both studies apply the AKM estimator, finding that firm effects explain 12 percent (Song et al., 2018) and 14 percent (Sorkin, 2018) of the variation in log earnings.⁸ By comparison, our AKM estimates suggest that firm effects explain 9 percent of the variation in log earnings. The key reason for this discrepancy seems to be the sample restrictions. We only include workers with earnings above the full-time minimum wage threshold. By comparison, Song et al. (2018) and Sorkin (2018) include individuals who work part time as long as their annual earnings exceed 25 percent of the full-time minimum wage threshold. In Online Appendix Figure B.2, we investigate what happens if we use their earnings cutoff. The AKM estimates then suggest that firm effects explain 11 percent of the variation of log earnings, which is nearly identical to the estimate of Song et al. (2018) and fairly similar to Sorkin (2018). This figure also shows estimates for a range of alternative earnings cutoffs. While the total log earnings variance increases substantially as the earnings cutoff decreases, the share of variation explained by AKM firm effects is relatively stable.

To investigate the second question of the sensitivity of the bias-corrected estimates, we restrict attention to workers and firms from a set of smaller states. This is necessary because it is computationally infeasible to apply Andrews et al. (2008) and Kline et al. (2018b) to the entire U.S. data. Online Appendix Figure B.5 compares the results from these alternative procedures for correcting for limited mobility bias to the estimates from BLM and AKM. The conclusion is clear: Limited mobility bias leads to upward bias in the AKM estimate of the firm component and downward bias in the AKM estimate of the worker sorting component. On average across

⁷Sorkin (2018) and Song et al. (2018) point out that limited mobility may bias their AKM estimates. In an attempt to investigate this issue, Sorkin (2018) also performs a few checks, including restricting the sample to large firms and splitting the sample in half on the basis of workers (which lets him compare results from two separate samples). Limited mobility bias, however, is about having few movers per firm, not small firms. Furthermore, the checks he perform involve significant changes in the composition of firms and workers in the estimation sample, in part because the connected set changes. Thus, it is not clear what, if any, conclusions one may draw about limited mobility bias from these checks.

⁸See Song et al. (2018), Appendix Table 2, and Sorkin (2018), Appendix Table 1.

Years:		2001-2008	2008-2015	Pooled
Panel A.		AK	M Estimatio	n
Share explained by:				
i) Worker Effects	$Var(x_i)$	75%	75%	75%
ii) Firm Effects	$Var(\psi_{i(i)})$	9%	9%	9%
iii) Sorting	$2Cov(x_i, \psi_{j(i)})$	5%	6%	5%
Sorting Correlation:	$Cor(x_i, \psi_{j(i)})$	0.09	0.11	0.10
Panel B.		BL	M Estimation	n
Share explained by:				
i) Worker Effects	$Var(x_i)$	72%	72%	72%
ii) Firm Effects	$Var(\psi_{i(i)})$	3%	3%	3%
iii) Sorting	$2Cov(x_i, \psi_{j(i)})$	13%	14%	14%
Sorting Correlation:	$Cor(x_i, \psi_{j(i)})$	0.43	0.46	0.44

Table 3: Decomposition results using AKM and BLM

Notes: This table presents the decomposition of log earnings variation using the AKM and BLM estimators for two time intervals. The analysis uses both workers who move between firms and stayers.

the states we consider, AKM suggests that firm effects explain more than 11 percent of the variation in log earnings. By contrast Andrews et al. (2008) and Kline et al. (2018b) suggest that firm effects explain about 5 percent, whereas the BLM method produces an estimate of firm effects around 2-3 percent.

Inequality within and between firms

We now shift attention to describing the inequality within and between firms. To do so, we follow Song et al. (2018) in expressing the variance of log earnings as:

$$Var(\log W_{it}) = \underbrace{Var(\log W_{it} - \mathbb{E}[\log W_{it}|j(i,t) = j])}_{Within-firm} + \underbrace{Var(\mathbb{E}[\log W_{it}|j(i,t) = j])}_{Between-firm}$$
(9)

$$= \underbrace{Var(x_i + \mathcal{X}'_{it}b - \mathbb{E}[x_i + \mathcal{X}'_{it}b|j(i,t) = j])}_{Worker heterogeneity within firms} + \underbrace{Var(\epsilon_{it})}_{Residual}$$
(9)

$$+ \underbrace{Var(\psi_{j(i,t)})}_{Firm effects} + \underbrace{2Cov(x_i + \mathcal{X}'_{it}b, \psi_{j(i,t)})}_{Sorting} + \underbrace{Var(\mathbb{E}[x_i + \mathcal{X}'_{it}b|j(i,t) = j])}_{Segregation}$$

where the first equality expresses the variance of log earnings in terms of inequality within and between firms, and the second equality decomposes these terms into economically interpretable subcomponents. Our interest is centered on the last three subcomponents, which capture distinct sources of inequality between firms: dispersion of firm pay premiums ("Firm effects"); sorting of high earning workers into high paying firms ("Sorting"); and worker segregation which reflects differences in the quality of the workforce across firms ("Segregation"). Both worker sorting and segregation reflect non-random allocation of workers to firms. However, sorting matters for aggregate inequality, whereas segregation does not. This is because an increase in segregation will be offset by a reduction in within-firm inequality. Thus, changes in segregation by itself does not affect earnings inequality; it does, however, matter for the relative importance of inequality within versus between firms.

To perform the decomposition in (9), we use exactly the same sample as in Table 3 which includes both workers who move between firms and stayers. The results are presented in Table 4. In Panel A, we report the terms in the first equality. We find that around one-third of the variance of log earnings can be accounted for by the dispersion of average earnings between firms. The remainder is due to heterogeneity across workers within firms. A comparison of the estimates across the two first columns suggests the between firm component has become slightly more important for inequality over time. This finding is broadly consistent with the results reported in Song et al. (2018).⁹

In the next two panels of Table 4, we use the procedures of AKM and BLM to estimate the subcomponents from the second equality. There are three main findings from this analysis. First, a vast majority of the inequality within firms can be accounted for by the observable characteristics and the fixed effects of the workers. Indeed, only 16 percent of the within-firm inequality reflects time-varying unobservables of the worker. Second, once one addresses limited mobility bias then firm effects explain only 10 percent of the inequality between firms. By comparison, sorting of high earning workers to high paying firms accounts for 40 percent while the remaining 50 percent can be attributed to worker segregation that is unrelated to firm pay premiums. Third, there seems to be little if any changes in the relative importance of firm effects, worker sorting and segregation over the time intervals we consider.

Our finding of the inequality contribution from firm effects changing little over time is consistent with Song et al. (2018), albeit their analysis uses AKM and thus suffers from limited mobility bias. Table 4 reveals, however, that this bias does not change materially over the time intervals we consider. As a result, bias correction seems to be empirically important for accurately describing the cross-sectional distribution of earnings in the U.S., but not for understanding the growth in earnings inequality.¹⁰

3.3.2 Extensions

The assumptions that $\phi_{ij} = x_i + \psi_j$ and $\gamma = \Upsilon = 0$ implies strong restrictions on the wage structure. The absence of interactions between worker and firm effects rules out strong complementaries in production, as in Shimer and Smith (2000) and Eeckhout and Kircher (2011). The

 $^{^{9}}$ The analyses in Song et al. (2018) is based on data from 1978 to 2013. Over this longer time period, they show that earnings inequality increased considerably, primarly due to a significant rise in the dispersion of average earnings across firms. During the period we consider, however, Song et al. (2018) also report a modest increase in earnings inequality, overall and between firms.

¹⁰Song et al. (2018) also argue that increases in sorting and segregation caused a large increase in between-firm inequality from 1981 to 2013. At first sight, it would seem like this is inconsistent with our findings. However, most of these increases happen before our data start. During the intervals since 2001 that we consider, Song et al. (2018) report modest increases in the contributions to between-firm inequality from sorting and segregation and a modest decrease from firm effects, consistent with our AKM estimates.

Years:		2001-2008	2008-2015	Pooled
Panel A.		Total	Decomposit	ion
Within Firm Share:	$Var(w_{it} - \mathbb{E}[w_{it} j])$	67%	64%	66%
Between Firm Share:	$Var(\mathbb{E}[w_{it} j])$	33%	36%	34%
Panel B.		AKM	Decomposit	ion
Shares of Within Firm Variance:				
Worker Heterogeneity:	$Var(x_i + X'_{it}b - \mathbb{E}[x_i + X'_{it}b j])$	84%	85%	84%
Residual:	$Var(\epsilon_{it})$	16%	15%	16%
Shares of Between Firm Variance:				
Firm Effects:	$Var(\psi_j)$	27%	25%	26%
Segregation:	$Var(\mathbb{E}[x_i + X'_{it}b j])$	58%	59%	59%
Sorting:	$2Cov(x_i + X'_{it}b, \psi_j)$	15%	16%	15%
Panel C.		BLM	Decomposit	ion
Shares of Within Firm Variance:				
Worker Heterogeneity:	$Var(x_i + X'_{it}b - \mathbb{E}[x_i + X'_{it}b j])$	83%	84%	84%
Residual:	$Var(\epsilon_{it})$	17%	16%	16%
Shares of Between Firm Variance:				
Firm Effects:	$Var(\psi_j)$	10%	10%	10%
Segregation:	$Var(\mathbb{E}[x_i + X'_{it}b j])$	50%	50%	50%
Sorting:	$2Cov(x_i + X'_{it}b, \psi_j)$	40%	40%	40%

Table 4: Decomposition of Inequality Within and Between Firms using AKM and BLM

Notes: This table presents the decomposition of log earnings variation within and between firms using the AKM and BLM estimators for two time intervals. The analysis uses both workers who move between firms and stayers.

assumption of no pass through of firm and market shocks is at odds with our data and a large body of evidence from many other developed countries. Thus, investigating these assumptions seems important to draw credible conclusions about the functioning of the U.S. labor market.

Non-additivity and complementarities

The assumption that $\phi_{ij} = x_i + \psi_j$ implies that all workers who move from firm j to j' will experience an earnings change of $\psi_{j'} - \psi_j$, no matter their quality x_i . An informal way to assess this log additive structure is to perform an event study of the earnings changes experienced by workers moving between different types of firms. Card et al. (2013b) and Card et al. (2018) use matched employer-employee data from Germany and Portugal to perform such event-study analyses of the earnings changes experienced by workers moving between different types of firms. In Online Appendix Figure B.3, we perform the same exercise, but this time for our U.S. data. This analysis uses the movers sample. As in Card et al. (2013b) and Card et al. (2018), we define firm groups based on the average pay of coworkers.

The results from the event study mirror those reported in Card et al. (2013b) and Card et al. (2018). Workers who move to firms with more highly-paid coworkers experience earnings raises, while those who move in the opposite direction experience earnings decreases of similar magnitude. Additionally, the gains and losses for movers in opposite directions between any two groups of firms are relatively symmetric. By comparison, earnings do not change materially

when workers move between firms with similarly paid coworkers. Another relevant finding from the event study is that the earnings profiles of the various groups are all relatively stable in the years before and after a job move. This lends support to Assumption 2, as it suggests that worker mobility does not seem to depend strongly on the trends in earnings beforehand or afterwards. Lastly, it is interesting to observe that the gains and losses for movers seem to be permanent. In contrast, in a large class of search models with job ladders, moves to firms that currently pay less is rationalized by arguing that these firms will pay more in the future.

Although the event study results are consistent with the log additive functional form, we cannot rule out interaction effects between worker and firm effects. Indeed, Bonhomme et al. (2019) point out that even if the functional form is non-additive, the gains and losses may look symmetric if workers making upward moves are of the same quality as those making downward moves. More generally, the degree of asymmetry one observes in the event study depends both on the magnitudes of any interaction effects and on the extent to which workers making upward moves differ in quality from those making downward moves. Thus, the event study analysis needs to be interpreted with caution.

To obtain an actual estimate of the importance of interactions between worker and firm effects, we follow BLM in using the following model of earnings:

$$w_{it} = \underbrace{\theta_{j(i,t)} \cdot x_i}_{\text{interaction}} + \psi_{j(i,t)} + \epsilon_{it} \tag{10}$$

which reduces to AKM when θ_j is the same for all firms. Under Assumptions 1 and 2 , we obtain:

$$\mathbb{E}[w_{it+1}|j_2 \to j_1] - \mathbb{E}[w_{it}|j_1 \to j_2] = \theta_{j_1} \left(\mathbb{E}\left[x_i|j_2 \to j_1\right] - \mathbb{E}\left[x_i|j_1 \to j_2\right]\right)$$
$$\mathbb{E}[w_{it+1}|j_1 \to j_2] - \mathbb{E}[w_{it}|j_2 \to j_1] = -\theta_{j_2} \left(\mathbb{E}\left[x_i|j_2 \to j_1\right] - \mathbb{E}\left[x_i|j_1 \to j_2\right]\right)$$

where $j_1 \rightarrow j_2$ $(j_2 \rightarrow j_1)$ is an indicator for a worker moving from firm 1 to 2 (firm 2 to 1). As long as the workers moving from 1 to 2 are not exactly the same as those moving from 1 to 2, the right hand side of these equalities are non-zero and we can recover θ_{j1}/θ_{j2} from the moment condition:

$$\frac{\mathbb{E}[w_{it+1}|j_2 \to j_1] - \mathbb{E}[w_{it}|j_1 \to j_2]}{\mathbb{E}[w_{it}|j_2 \to j_1] - \mathbb{E}[w_{it+1}|j_1 \to j_2]} = \frac{\theta_{j1}}{\theta_{j2}}$$
(11)

Thus, provided that the composition of movers differs across firms, it is possible to identify θ_j (up to scale) for every firm. To take (10) to the data, however, it is useful to reduce the number of parameters to estimate. As above, we follow BLM in classifying firms to ten types according to the empirical earnings distribution within firms. Then we restrict θ_j to be the same for all firms of a given type. We refer to Section 5.4 for details on the estimation procedure.

Figure 3 displays the estimated nonlinearities. We plot the means of log earnings for each firm type and at 10 deciles of worker heterogeneity. On the x-axis, firm types are ordered in ascending order, where "lower" and "higher" types refer to low and high mean log earnings. The results show clear evidence of worker heterogeneity: For the same type of firm, better workers



Figure 3: Estimates of Interactions between Firm and Worker Effects using BLM

Notes: In this figure, we present estimates of interactions between firm and worker effects using the BLM estimator. We plot the means of log earnings for each firm type and deciles of worker heterogeneity. On the x-axis, firm types are ordered in ascending order, where "lower" and "higher" types refer to low and high mean log earnings.

earn significantly more. For a given worker, there is also some variation in log earnings between firm types, although to a lesser extent. As shown in equation (11), the parameters governing nonlinearities are identified from comparing the gains from moving from a low to a high type of firm for workers of different quality. As evident from Figure 3, the gains from such a move are considerably larger for better workers. For example, moving from the lowest to the highest type of firm increases earnings by 22, 47 and 78 percentage points for individuals at the 20, 50 and 80 percentile in the worker quality distribution.

The evidence of nonlinearities raises several questions. To what extent do interaction effects bias the estimates from the log additive model? Are nonlinearities empirically important as a source of earnings inequality? In Table 5, we investigate these questions by extending the AKM decomposition to incorporate the contribution from interactions between worker and firm effects. Re-arranging equation (10), we get

$$w_{it} = \underbrace{\bar{\theta}(x_i - \bar{x})}_{\tilde{x}_i} + \underbrace{\left(\psi_{j(i,t)} + \theta_{j(i,t)}\bar{x}\right)}_{\tilde{\psi}_{j(i,t)}} + \underbrace{\left(\theta_{j(i,t)} - \bar{\theta}\right)(x_i - \bar{x})}_{\varrho_{ij(i,t)}} + \epsilon_{it} \tag{12}$$

where $\bar{\theta} \equiv \mathbb{E}\left[\theta_{j(i,t)}\right]$ and $\bar{x} \equiv \mathbb{E}\left[x_i\right]$. This equation decomposes the earnings of worker *i* in period *t* into three distinct components: \tilde{x}_i gives the direct effect of the quality of worker *i* (evaluated at the average firm), $\tilde{\psi}_{j(i,t)}$ represents the direct effect of firm *j* (evaluated at the

		Μ	odel Sp	ecificati	on
		(1)	(2)	(3)	(4)
Share explained by:					
i) Worker Quality	$Var(x_i)$	72.4%	70.4%	73.5%	71.6%
ii) Firm Effects	$Var(\psi_{i(i)})$	3.2%	4.3%	3.0%	4.3%
iii) Sorting	$2Cov(x_i, \psi_{j(i)})$	12.9%	13.1%	12.8%	13.1%
iv) Interactions	$Var(\rho_{ij})$		3.0%		3.3%
	$+2Cov(x_i+\psi_{j(i)},\varrho_{ij})$		-1.8%		-2.5%
v) Time-varying Effects	$Var(\psi_{j(i),t} - \psi_{j(i)}) + 2Cov(x_i, \psi_{j(i),t} - \psi_{j(i)})$			0.3%	0.3%
Sorting Correlation:	$Cor(x_i, \psi_{i(i)})$	0.43	0.38	0.43	0.37
Variance Explained:	R^2	0.89	0.89	0.90	0.90
Specification:					
Firm-Worker Interactions		X	\checkmark	X	\checkmark
Time-varying Firm Effects		×	×	\checkmark	\checkmark

Table 5: Comparison of BLM Specifications

Notes: This table presents the decomposition of log earnings variation into firm and worker effects using the BLM estimator for four specifications: baseline, allowing for worker effects to interact with firm effects ("Firm-Worker Interactions"), allowing for a time-varying component in the firm effects due to the pass through of value added shocks ("Time-varying Firm Effects"), and allowing for both interactions between firm and worker effects and time-varying firm effects. The analysis uses both workers who move between firms and stayers.

average worker), and $\rho_{ij(i,t)}$ captures the interaction effect between firm j and worker i quality. Using equation (12), we obtain a new variance decomposition of log earnings:

$$Var(w_{it}) = Var\left[\tilde{x}_{i}\right] + Var\left[\tilde{\psi}_{j(i,t)}\right] + 2Cov\left[\tilde{x}_{i}, \tilde{\psi}_{j(i,t)}\right] + Var\left[\varrho_{ij(i,t)}\right] + 2Cov\left[\tilde{x}_{i} + \tilde{\psi}_{j(i,t)}, \varrho_{ij(i,t)}\right]$$
(13)

The first three components are informative about the inequality contribution from worker effects, firm effects and worker sorting, net of interaction effects. The last two components are informative about the inequality contribution from interaction effects, as measured by the dispersion of $\rho_{ij(i,t)}$ across firms and the extent to $\rho_{ij(i,t)}$ is larger in firms with high wages. If $\theta_j = \bar{\theta}$ for every firm j, then these two components would be zero, and the decomposition in (13) reduces to the standard AKM decomposition.

The results from the decomposition in (13) are presented in column (2) of Table 5. Our estimates suggest the dispersion of interaction effects across firms explains three percent of the earnings inequality. However, the total contribution to earnings inequality from nonlinearities is muted by the interaction effects being larger in firms with higher paid workers. We also find that omitting interaction effects causes a downward bias in the firm effects and an upward bias in the worker effects.

Pass through of shocks and time-varying types

The assumption that $\gamma = \Upsilon = 0$ restricts firm effects to be constant over time. However, the significant pass-through rates imply that firm effects actually evolve over time as employers experience changes in the value added at the firm or market level. To capture this, we now let γ differ from Υ and propose an adjustment to the AKM model which allows us to isolate the time-invariant component of the firm effects.

Our approach proceeds in two steps. First, we construct an adjusted earnings measure by removing the time-varying firm and market specific component of earnings. To do so, we use the firm and market level value added multiplied by the estimated passthrough coefficients at the firm and market level (see Table 2). Second, we recover the time-invariant firm and worker effects by applying the methods of AKM or BLM to the adjusted measure of earnings. Consider the following adjusted two-way specification for earnings of workers across firms:

$$\mathbb{E}[w_{it} - \gamma(y_{j(i,t),t} - y_{j(i,t),1}) - (\Upsilon - \gamma) \left(\bar{y}_{r(i,t),t} - \bar{y}_{r(i,t),1}\right) | j(i,1), ..., j(i,T)] = x_i + \psi_j.$$

The left-hand side removes the earnings dynamics due to passthrough of firm-specific shocks, $\gamma(y_{j(i,t),t} - \bar{y}_{r(i,t),t})$, and market shocks, $\gamma \bar{y}_{r(i,t),t}$. What remains is the worker effect x_i and the time-invariant firm effect ψ_j , which can be estimated by applying AKM or BLM to the adjusted earnings measure.

In column (3) of Table 5, we extend the BLM decomposition of the variance of log earnings in (9) to incorporate the contribution from time-invariant and time-varying firm effects. We find that time-varying firm effects explain little if any of the variation in log earnings, and that the importance of firm effects and worker sorting do not change materially if we take the pass through of firm shocks into account. Comparing the results in column (4) to those presented in column (2) shows that time variation also has little to no explanatory power when accounting for nonlinearities.

3.4 How the empirical findings inform our modeling choices

Before we present the model of the labor market, it is useful to explain how our modeling choices are guided by the empirical findings reported above.

One of these findings is that most of the dispersion in firm average earnings is driven by heterogeneity in the composition of workers who are employed at different firms. This finding highlights the importance of allowing for heterogeneity in worker productivity and sorting in the labor market. To accommodate this, we develop an equilibrium model of the labor market with two-sided heterogeneity.

A related finding is the existence of firm wage premiums, although they explain relatively little of the earnings inequality once one corrects for limited mobility bias. In the textbook model of a perfectly competitive labor market, the law of one price holds and there should exist a single market wage for a given quality of a worker, no matter the employer for which she works. Firm wage premiums are at odds with this theory. To allow for firm effects, we let employers differ in workplace characteristics or amenities. Employers that offer favorable working conditions attract labor at lower than average wages, whereas employers offering unfavorable working conditions must pay premiums as offsetting compensation in order to attract labor.

Another important finding is that idiosyncratic value added shocks to a firm transmit to the earnings of incumbent workers. To capture such a pass through, we will allow workers to have heterogeneous preferences over the same amenities. As a result of this preference heterogeneity, the employer faces an upward sloping supply curve for labor, implying that wage is an increasing function of size. We will assume that employers do not observe the idiosyncratic taste for amenities of any given worker. This information asymmetry implies that employers cannot price discriminate with respect to workers' reservation values. Instead, if a firm becomes more productive and thus wants to increase its size, the employer must offer higher wages to all workers.

A related finding is that the earnings of incumbent workers respond significantly less to an idiosyncratic value added shock to the current firm than to a (same size) shock to all firms in their current industry and area. One way to rationalize this finding is to partition the worker's choice set into "nests" by the industry and location of the firms, and allow that firms in the same nest are closer substitutes in terms of their amenities than firms in other nests. This nested structure on the workers' preferences captures that many job characteristics are likely to vary systematically both across areas and industries.

Another finding is that large differences in value added across firms do not create sizable firm wage premiums. This finding would be expected if employers face perfectly elastic supply curves for labor; a productive firm could then increase its size without raising wages. However, given the evidence of upward sloping labor supply curves, one might expect sizable wage premiums from being employed in more productive firms. To rationalize that the observed differences in value added across firms do not create larger firm wage premiums, we allow for correlation between the firms' productivity and amenities. Such a correlation could, for example, arise because productive firms invest in better amenities as a way to attract better workers, which pushes down the wages that these firms have to pay and thus attenuates the wage premiums.

Lastly, the gains in log earnings from moving to higher paying firms are considerably larger for better workers. These positive interaction effects are especially pronounced at the upper tails of the distributions of worker and firm effects. To capture such interaction effects, we allow for (but do not impose) production complementarities in the firm's technology. Such complementarities could help explain the pattern of sorting we observe in the data as they incentivize better workers to sort into more productive firms, as in Shimer and Smith (2000) and Eeckhout and Kircher (2011). A competing explanation for the sorting patterns would be that high ability workers differ in how much they value amenities as compared to low ability workers. By taking our model to the data, it is possible to quantify the strength and implications of worker sorting based on workplace amenities versus production complementarities.

4 Model of the labor market

This section develops a model of the labor market where multiple employers compete with one another for workers who have heterogeneous preferences over non-wage job characteristics or amenities.

4.1 Agents, preferences and technology

The economy is composed of a large number of workers indexed by i and a large set of firms indexed by j = 1, ..., J. Each firm belongs to a market r(j). Let J_r denote the set of firms in market r. We will rely on the approximation that firms employ many workers and that each market has many firms. For tractability, we assume that workers, firms and markets face exogenous birth-death processes which ensure stationarity in the productivity distributions of workers, firms and markets.

Worker productivity and preferences

Workers are heterogeneous both in preferences and productivity. The productivity of worker i is given by (X_i, V_{it}) where X_i is the permanent component and V_{it} is the time-varying component. In period t, worker i has the following preferences over alternative firms j and earnings W:

$$u_{it}(j,W) = \log \tau W^{\lambda} + \log G_j(X_i) + \beta^{-1} \epsilon_{ijt}$$
(14)

where $G_j(X)$ denotes the value that workers of quality X are expected to get from the amenities that firm j offers, and ϵ_{ijt} denotes worker i's idiosyncratic taste for the amenities of firm j. The parameters (τ, λ) describe the tax function that maps wages to income available for consumption. In Section 5, we show how well this parsimonious tax function approximates the tax rates of the US tax system.

The specification of preferences in equation (14) allows for the possibility that workers view firms as imperfect substitutes. Fixing the permanent component of worker productivity X, the preference term $G_j(X)$ gives rise to *vertical* employer differentiation: some employers offer good amenities while other employers have bad amenities. Our preference specification combines this vertical differentiation with *horizontal* employer differentiation: workers are heterogeneous in their preferences over the same firm. This horizontal differentiation has two distinct sources. The first is that $G_j(X)$ varies freely across values of X. Thus, we permit systematic heterogeneity in the preferences for a given firm depending on the permanent component of worker productivity. The second is the idiosyncratic taste component ϵ_{ijt} . The importance of this second source of horizontal differentiation is governed by the parameter β , which tells us the variability across workers in the idiosyncratic taste for a given firm. Formally, this parameter is proportional to the inverse of the standard deviation of ϵ_{ijt} in log-dollars.

We assume that $(\epsilon_{i1t}, ..., \epsilon_{iJt}) \equiv \vec{\epsilon}_{it} \sim \Psi(\vec{\epsilon} | \vec{\epsilon}_{it-1}, X_i)$ follows a Markov process with independent innovations across individuals with the same productivity X. This assumption does not

imply strong restrictions on the copula (and, by extension, the patterns of mobility by worker quality) over time. We assume, however, that the (cross-sectional) distribution of $\vec{\epsilon}_{it}$ has a nested logit structure in each period:

$$F(\vec{\epsilon}_{it}) = \exp\left[-\sum_{r} \left(\sum_{j \in J_r} e^{-\frac{\epsilon_{ijt}}{\rho_r}}\right)^{\rho_r}\right].$$

This structure allows the preferences of a given worker to be correlated across alternatives within each nest. In the empirical analysis, we specify the nest as the combination of industry and area, and refer to it as a market. The parameter ρ_r measures of the degree of independence in a worker's taste for the alternative firms within market r, i.e. $\rho_r = \sqrt{1 - corr(\epsilon_{ijt}, \epsilon_{ij't})}$ if r(j) = r(j') = r. Thus, this parameter is equal to zero if each worker views firms within the same market as perfect substitutes, while it is equal to one if the worker views these firms as completely independent alternatives.

Firm productivity and technology

Importantly, we let firms differ not only in workplace amenities but also in terms of productivity and technology. We start by introducing the total efficiency unit of labor at the firm:

$$L_{jt} = \iint V \cdot X^{\theta_j} \cdot D_{jt}(X, V) \mathrm{d}X \mathrm{d}V,$$

where θ_j is a firm-specific parameter that determines the returns to labor permanent productivity X, and $D_{jt}(X, V)$ is the mass of workers with productivity (X, V) demanded by the firm. The revenues Y_{jt} generated by firm j in period t is determined by the production function

$$Y_{jt} = A_{jt} L_{jt}^{1 - \alpha_{r(j)}}$$

where A_{jt} is the firm's productivity (TFP), and $\alpha_{r(j)}$ is the market-specific scale parameter. This specification of the revenue production function abstracts from capital, or equivalently, assumes that capital can be rented at some fixed price. However, the specification does not require the product market to be competitive. As shown in Online Appendix C.5, it is possible to derive the same specification of the revenue production function (and, by extension, labor demand) if firms have price-setting power in the product market.

It is useful to express the productivity component A_{jt} as:

$$\begin{aligned} A_{jt} &= \bar{A}_{r(j)t}\tilde{A}_{jt} \\ &= \bar{P}_{r(j)}\bar{Z}_{r(j)t}\tilde{P}_{j}\tilde{Z}_{j} \end{aligned}$$

where $\bar{A}_{r(j)t}$, $\bar{P}_{r(j)}$, and $\bar{Z}_{r(j)t}$ represent the overall, the permanent and the time-varying components of productivity that are shared by all firms in market r, while \tilde{A}_{jt} , \tilde{P}_j and \tilde{Z}_{jt} denote the overall, the permanent and the time-varying components that are specific to firm j. Let $W_{jt}(X, V)$ denote the wage that firm j offers to workers with characteristics (X, V) in period t and $B_{jt} = \iint W_{jt}(X, V) D_{jt}(X, V) dX dV$ denote the wage bill of the firm, i.e. the total sum of wages paid to its workers. The profit of the firm is then given by $\Pi_{jt} = Y_{jt} - B_{jt}$.

Importantly, we do not restrict the relationship between amenities $G_j(X)$, permanent productivity components $(\bar{P}_{r(j)}, \tilde{P}_j)$, and technology $(\theta_j, \alpha_{r(j)})$. As a result, there are several reasons why better workers may prefer to work in productive firms. One possible reason is that productive firms (high A) could have better amenities (high G) and high wage workers (high X) may value amenities more than low wage workers. Another possible reason is strong complementarities in production, as in Shimer and Smith (2000) and Eeckhout and Kircher (2011). By letting θ_j vary across firms, the specification of technology permits such complementarities between productive (high X) workers and productive (high A) firms.

4.2 Information, wages and equilibrium

We consider an environment where all labor is hired in a spot market and ϵ_{ijt} is private information to the worker. Hence, the wage may depend on the worker's quality (X, V), but not her value of ϵ_{ijt} . Given the set of offered wages $\mathbf{W}_t = \{W_{jt}(X, V)\}_{j=1,...,J}$ by all firms, worker *i* chooses a firm *j* to maximize his utility u_{it} in each period:

$$j(i,t) \equiv \arg\max_{j} u_{it}(j, W_{jt}(X_i, V_{it})).$$
(15)

We introduce a wage index at the level of the market r defined by:

$$I_{rt}(X,V) \equiv \left(\sum_{j'\in J_r} \left(\tau^{1/\lambda} G_{j'}(X)^{1/\lambda} W_{j't}(X,V)\right)^{\frac{\lambda\beta}{\rho_r}}\right)^{\frac{\rho_r}{\lambda\beta}}$$
(16)

from which we can derive the probability that an individual of type (X, V) chooses to work at firm j given all offered wages in the economy:

$$Pr[j(i,t)=j|X_i=X, V_{it}=V, \mathbf{W}_t] = \frac{\left(I_{r(j)t}(X,V)\right)^{\lambda\beta}}{\sum_{r'} I_{r't}(X,V)^{\lambda\beta}} \left(\tau^{1/\lambda}G_j(X)^{1/\lambda}\frac{W_{jt}(X,V)}{I_{r(j)t}(X,V)}\right)^{\frac{\lambda\beta}{\rho_{r(j)}}}$$

We consider an equilibrium where the firm views itself as infinitesimal within the market. Thus, given the total mass of workers N and the stationary cross-sectional distributions of Xand V, $M_X(X)$ and $M_V(V)$, employer j considers the following firm-specific labor supply curve when setting wages $W_{jt}(X, V)$:

$$S_{jt}(X,V,W) \equiv NM_X(X)M_V(V)\frac{\left(I_{r(j)t}(X,V)\right)^{\lambda\beta}}{\sum_{r'}I_{r't}(X,V)^{\lambda\beta}} \left(\tau^{1/\lambda}G_j(X)^{1/\lambda}\frac{W}{I_{r(j)t}(X,V)}\right)^{\frac{\lambda\beta}{p_{r(j)}}}.$$

This means the firm ignores the negligible effect of changing its own wages on the market level wage index $I_{rt}(X, V)$. Then each firm chooses labor demand $D_{jt}(X, V)$ by setting wages $W_{jt}(X, V)$ for every type of worker (X, V) to maximize profits subject to the labor supply constraint $S_{jt}(X, V, W)$. For each firm j and time t,

$$\Pi_{jt} = \max_{\{W_{jt}(X,V)\}_{(X,V)}} A_{jt} \left(\iint X^{\theta_j} V D_{jt}(X,V) dX dV \right)^{1-\alpha_{r(j)}} - \iint W_{jt}(X,V) D_{jt}(X,V) dX dV$$

s.t. $D_{jt}(X,V) = S_{jt}(X,V,W_{jt}(X,V))$ for each X,V (17)

Under these simplifying assumptions, the definition of the equilibrium naturally follows from the presented environment:

Definition 1. Given firm characteristics $(\alpha_{r(j)}, A_{jt}, \theta_j)_{j,t}$, worker distributions $N, M_X(\cdot), M_V(\cdot)$, preference parameters $(\beta, \rho_r, G_j(\cdot))$ and tax parameters (λ, τ) , we define the equilibrium as the worker decisions j(i, t), market level wage indices $I_{rt}(X, V)$, firm-specific labor supply curves $S_{jt}(X, V, W)$, wages $W_{jt}(X, V)$ and labor demand $D_{jt}(X, V)$ such that:

- i. Workers choose firms that maximize their utility, as defined in equation (15).
- ii. Firms choose labor demand $D_{jt}(X, V)$ by setting wages $W_{jt}(X, V)$ for every type of worker (X, V) to maximize profits subject to the labor supply constraint $S_{jt}(X, V, W)$, as described in equation (17).
- iii. The market level constants $I_{rt}(X, V)$ are generated from the workers' optimal decisions j(i, t).

In Lemma 7 in Online Appendix C.1, we show the uniqueness of the equilibrium which proves useful in the estimation of the model and is needed for the counterfactual analysis.

4.3 Structural equations

As shown in Proposition 1 in Online Appendix C.1, our model delivers the following structural equations for wages, value added and wage bills for firm $j \in J_r$:

$$w_j(x, v, \bar{a}, \tilde{a}) = \theta_j x + v + c_r - \alpha_r h_j + \frac{1}{1 + \alpha_r \lambda \beta} \bar{a} + \frac{1}{1 + \alpha_r \lambda \beta / \rho_r} \tilde{a}$$
(18)

$$y_j(\bar{a},\tilde{a}) = (1 - \alpha_r)h_j + \frac{1 + \lambda\beta}{(1 + \alpha_r\lambda\beta)}\bar{a} + \frac{1 + \lambda\beta/\rho_r}{1 + \alpha_r\lambda\beta/\rho_r}\tilde{a}$$
(19)

$$b_j(\bar{a},\tilde{a}) = c_r + (1 - \alpha_r)h_j + \frac{1 + \lambda\beta}{(1 + \alpha_r\lambda\beta)}\bar{a} + \frac{1 + \lambda\beta/\rho_r}{1 + \alpha_r\lambda\beta/\rho_r}\tilde{a}$$
(20)

where we use lower case letters to denote logs (e.g. $x \equiv \log X$) and c_r is a market-specific constant that is equal to $\log \frac{(1-\alpha_r)\lambda\beta/\rho_r}{1+\lambda\beta/\rho_r}$. These equations describe how the *potential* outcomes of workers and firms are determined, that is, they tell us the realization of $w_{jt}(x,v)$, b_{jt} , and y_{jt} that would have been experienced had worker productivity (x,v), firm TFP \tilde{a} and market TFP \bar{a} been exogenously set.

The equations in (18)-(20) show that $w_j(x, v, \bar{a}, \tilde{a})$, $b_j(\bar{a}, \tilde{a})$, and $y_j(\bar{a}, \tilde{a})$ depend on the same four components: the constant c_r , the component of productivity that is specific to the firm \tilde{a} , the component of productivity that is common to firms in the same market \bar{a} , and an amenity component h_j . Comparing the wage equation to the other structural equations, we see that $w_j(x, v, \bar{a}, \tilde{a})$ also depends on the worker's own productivity (x, v). Moreover, workers with the same productivity are paid differently depending on the firm-specific parameter θ_j . As expected, if a firm j becomes more productivity (\tilde{a} or \bar{a} increase) then $y_j(\bar{a}, \tilde{a})$ increases. Because firm j has become more productive, it will demand more labor, which raises $w_j(x, v, \bar{a}, \tilde{a}), b_j(\bar{a}, \tilde{a})$, and $y_j(\bar{a}, \tilde{a})$.

From the structural wage equations, we also obtain:

$$\ell_j(\bar{a},\tilde{a}) = h_j + \frac{\lambda\beta}{(1+\alpha_r\lambda\beta)}\bar{a} + \frac{\lambda\beta/\rho_r}{1+\alpha_r\lambda\beta/\rho_r}\tilde{a}$$
(21)

where h_j can be interpreted as the efficiency unit of labor the firm would have if \tilde{a} and \bar{a} are exogenously set to zero. This "TFP-neutral" notion of the total efficiency of labor depends on the quality and quantify of the workforce. We can then write

$$h_j = \underbrace{\log \mathbb{E}\left[X_i^{\theta_{j(i,t)}} \mid j(i,t)=j\right]}_{\bar{x}_j} + \bar{g}_j.$$

The component \bar{g}_j is a weighted average of the X-specific amenities of firm j. It captures the vertical differentiation of firms. All else equal, a greater \bar{g}_j raises the size of the firm, thus increasing its wage bill and value added. At the same time, better amenities push down the wages that these firms have to pay their workers, lowering $w_j(x, v)$. The other component of h_j is the average quality of the workers in the firm \bar{x}_j , reflecting sorting of workers to firms. If $G_j(X)$ increases by the same rate for each X in a given firm j, then \bar{g}_j increases while \bar{x}_j is unaffected. Hence, \bar{x}_j captures that some firms may have relatively good amenities for some worker types. All else equal, better co-workers raise the total labor efficiency of the firm's workforce. With decreasing return to scale in the production function ($\alpha > 0$), this also lowers the wages that firm j pays to a given worker.

Another important feature of the structural equations is that they are log additive in the arguments $\theta_j x$, $\bar{a}_{r(j)t}$, \tilde{a}_{jt} , and h_j . This log additivity is useful for several reasons. First, it makes it straightforward to quantify the relative importance of the determinants of worker and firm outcomes. Second, it forges a direct link between the structural wage equation and the two-way fixed effect models discussed in Section 3.3. This link will help interpret the variance decomposition of AKM and BLM through the lens of the model. Third, it makes it possible to prove identification of the parameters of the model, as shown in Section 5.

4.4 Rents, compensating differentials, and allocative inefficiencies

We conclude the presentation of the model by defining the rents that workers and employers earn from ongoing employment relationships and explaining how these quantities relate to reservation wages, compensating wage differentials, and allocative inefficiencies.

Worker rents

In our model, rents are due to the idiosyncratic taste component ϵ_{ijt} which gives rise to horizontal differentiation of firms, upward sloping labor supply curves, and employer wage setting power. We assumed that employers do not observe the idiosyncratic taste for amenities of any given worker. This information asymmetry implies that firms cannot price-discriminate with respect to workers' reservation wages. As a result, the equilibrium allocation of workers to firms creates surpluses or rents for inframarginal workers, defined as the excess return over that required to change a decision, as in Rosen (1986). In our model, worker rents may exist at both the firm and the market level:

Result 1. We define the *firm level rents* of worker i, R_{it}^w , as the surplus he derives from being inframarginal at his current choice of firm. Given his equilibrium choice j(i, t), R_{it}^w is implicitly defined by:

$$u_{it}(j(i,t), W_{j(i,t),t}(X_i, V_{it}) - R_{it}^{w}) = \max_{j' \neq j(i,t)} u_{it}(j', W_{j',t}(X_i, V_{it})).$$

As shown in Lemma 9 in the Online Appendix, the expected worker rents at the firm level is given by:

$$\mathbb{E}[R_{it}^{w}|j(i,t)=j] = \frac{1}{1+\lambda\beta/\rho_{r(j)}}\mathbb{E}[W_{jt}(X_i,V_{it})|j(i,t)=j].$$

Result 2. We define the **market level rents** of worker *i*, R_{it}^{wm} , as the surplus derived from being inframarginal at his current choice of market. Given his equilibrium choice r(j(i,t)), R_{it}^{wm} is implicitly defined by:

$$u_{it}(j(i,t), W_{j(i,t),t}(X_i, V_{it}) - R_{it}^{wm}) = \max_{j' \mid r(j') \neq r(j(i,t))} u_{it}(j', W_{j',t}(X_i, V_{it}))$$

As shown in Lemma 9 in the Online Appendix, the expected worker rents at the market level is given by:

$$\mathbb{E}\left[R_{it}^{wm}|j(i,t)=j\right] = \frac{1}{1+\lambda\beta}\mathbb{E}\left[W_{jt}(X_i, V_{it})|j(i,t)=j\right]$$

Market level rents exceed firm level rents whenever the next best firm is in the same market as the current choice of firm. If the preferences of a given worker are independent across firms within each market, then the next best firm will almost surely be in a different market. If, on the other hand, these preferences are correlated then there could well exist other firms within the same market that are close substitutes to the current firm. The next best firm may then very well be in the same market as the current choice of firm, in which case R_{it}^{wm} will exceed R_{it}^{w} .

To interpret the measure of firm level rents and link it to compensating differentials, it is useful to express R_{it}^{w} in terms of reservation wages. The worker's reservation wage for his current choice of firm is defined as the lowest wage at which he would be willing to continue working in this firm. Substituting in preferences in the above definition of R_{it}^w , we get:

$$\underbrace{\log W_{j(i,t),t}(X_i, V_{it})}_{\text{current wage}} - \underbrace{\log \left(W_{j(i,t),t}(X_i, V_{it}) - R_{it}^w \right)}_{\text{reservation wage}} = \underbrace{\log W_{j(i,t),t}(X_i, V_{it})}_{\text{current wage}} - \underbrace{\log W_{j^o(i,t),t}(X_i, V_{it})}_{\text{wage at best outside option}} + \underbrace{\log G_{j(i,t)}^{1/\lambda}(X_i)e^{\frac{1}{\lambda\beta}\epsilon_{ij^o(i,t)}t}}_{\text{current amenities}} - \underbrace{\log G_{j^o(i,t)}^{1/\lambda}(X_i)e^{\frac{1}{\lambda\beta}\epsilon_{ij^o(i,t)}t}}_{\text{amenities at best outside option}}$$
(22)

The average worker choosing firm j may be far from the margin of indifference and would maintain the same choice even if his current firm offered significantly lower wages. The same thing holds true for the average worker choosing any other firm. The difference between the reservation wage and the actual wage is the rent earned by a person at his current choice of firm.

Compensating differentials

By definition, the marginal workers are indifferent between the current choice of firm and the next best option. They earn no rents as their reservation wages equal the actual wages paid by their current firms. By solving equation (22) for workers for whom reservation wages are equal to current wages, it becomes clear that the market wage differences are informative about the preferences for amenities of these marginal workers. Another way of saying this is that the equilibrium allocation of workers to firms is such that utility gains (or losses) of marginal workers due to the amenities of their firms are exactly offset by market wage differences. Thus, the market wage differences across firms define the equalizing or compensating wage differentials:

Result 3. The compensating differential between firm j and firm j' for workers of type (X, V) is defined as

$$CD_{jj't}(X,V) = u_{it}(j, W_{jt}(X,V)) - u_{it}(j', W_{jt}(X,V)) \text{ s.t. } R^w_{it} = 0$$

= $\log W_{jt}(X,V) - \log W_{j't}(X,V)$
= $(\theta_j - \theta_{j'})x + \psi_{jt} - \psi_{j't}$

where the last equality follows directly from equation (18) and

$$\psi_{jt} \equiv c_r - \alpha_r h_j + \frac{1}{1 + \alpha_r \lambda \beta} \bar{a}_{r(j),t} + \frac{1}{1 + \alpha_r \lambda \beta / \rho_r} \tilde{a}_{jt}$$

For any two firms j and j', there exists a distribution of compensating differentials. This distribution arises not because of heterogeneous tastes of workers, but because of differences in technology across firms. If θ_j does not vary across firms, there is only one compensating differential per employer. It is given by the firm effect ψ_{jt} , which is paid to all workers independent of their productivity.

Employer rents

The equilibrium allocation of workers to firms may also create surpluses or rents for employers. The employer rents arise because of the additional profit the firm can extract by taking advantage of its wage setting power. To measure employer rents, we therefore compare the profit Π_{jt} the firm actually earns to what it would have earned if the employer solved the firm's problem under the assumption that the labor supply it faces is perfectly elastic. In other words, wages, profits and employment are such that $D_{jt}^{pt}(X, V)$ solves the firm's profit maximization taking $W_{jt}^{pt}(X, V)$ as given:

$$\Pi_{jt}^{\text{pt}} = \max_{\left\{D_{jt}^{\text{pt}}(X,V)\right\}_{(X,V)}} A_{jt} \left(\iint VX^{\theta_j} \cdot D_{jt}^{\text{pt}}(X,V) \mathrm{d}X \mathrm{d}V\right)^{1-\alpha_{r(j)}} - \iint D_{jt}^{\text{pt}}(X,V) \cdot W_{jt}^{\text{pt}}(X,V) \mathrm{d}X \mathrm{d}V,$$

where the labor demand and the wages also have to clear:

$$S_{jt}(X, V, W_{jt}^{\text{pt}}(X, V)) = D_{jt}^{\text{pt}}(X, V) \text{ for all } t, j, X, V.$$

The only difference in this counterfactual environment is that the firm does not take into account its wage setting power through the upward-sloping labor supply curve. In other words, the firm behaves as if it faces a perfectly elastic labor supply curve, i.e. as if it was a "price taker"; thus the superscript pt. Similarly we define $W_{jt}^{\text{ptm}}(X, V)$, Π_{jt}^{ptm} and $D_{jt}^{\text{ptm}}(X, V)$ as the equilibrium outcome when all firms in a market act as price takers.

Result 4. We define the **employer rents** at the firm level R_{jt}^f and at market level R_{jt}^{fm} as the additional profit firm j derives from the presence of inframarginal workers. Given our previous definitions:

$$R_{jt}^{f} = \Pi_{jt} - \Pi_{jt}^{pt}$$

$$= \left(1 - \frac{\alpha_{r(j)} \left(1 + \lambda \beta / \rho_{r(j)}\right)}{1 + \alpha_{r(j)} \lambda \beta / \rho_{r(j)}} \left(\frac{\lambda \beta / \rho_{r(j)}}{1 + \lambda \beta / \rho_{r(j)}}\right)^{-\frac{(1 - \alpha_{r(j)}) \lambda \beta / \rho_{r(j)}}{1 + \alpha_{r(j)} \lambda \beta / \rho_{r(j)}}}\right) \Pi_{jt}$$

$$R_{it}^{fm} = \Pi_{it} - \Pi_{it}^{ptm}$$

$$(23)$$

$$= \left(1 - \frac{\alpha_{r(j)} \left(1 + \lambda\beta/\rho_{r(j)}\right)}{1 + \alpha_{r(j)}\lambda\beta/\rho_{r(j)}} \left(\frac{\lambda\beta/\rho_{r(j)}}{1 + \lambda\beta/\rho_{r(j)}}\right)^{-\frac{(1 - \alpha_{r(j)})\lambda\beta}{1 + \alpha_{r(j)}\lambda\beta}}\right) \Pi_{jt}$$
(24)

where the last equality in (23) and in (24) are shown in Lemma 10 and 11 in the Online Appendix.

It is important to observe that R_{jt}^f and R_{jt}^{fm} do not necessarily represent ex-ante rents. Suppose, for example, that each employer initially chooses the amenities offered to the workers by deciding on the firm's location, the working conditions, or both. Next, the employers compete with one another for the workers who have heterogeneous preferences over the chosen amenities. These heterogeneous preferences give rise to wage setting power which employers can use to extract additional profits or rents. Of course, the existence of such ex-post rents could simply be returns to costly choices of amenities.

Empirically, it is difficult to credibly distinguish between ex-ante and ex-post employer rents. It would require information (or assumptions) about how firms choose and pay for the amenities offered to workers. Given our data, we are severely limited in the ability to distinguish between ex-ante and ex-post rents. Instead, we assume that firms are endowed with a fixed set of amenities, or, more precisely, we restrict amenities to be fixed over the estimation window. It is important to note what is *not* being restricted under this assumption. First, it does not restrict whether or how amenities $G_j(X)$ relate to the technology parameters α_r , θ_j or the productivity components $\bar{P}_{r(j)}$ and \tilde{P}_j . Second, it neither imposes nor precludes that employers initially choose amenities to maximize profits. Indeed, it is straightforward to show that permitting firms to initially choose amenities would not affect any of our estimates. Nor would it matter for the interpretation of any result other than whether R_{jt}^f and R_{jt}^{fm} should be viewed as ex-ante or ex-post rents.

Wedges and allocative inefficiencies

We conclude the model section by investigating the questions of whether and in what situations the equilibrium allocation of workers to firms will be inefficient. We here present the key results, and refer to Online Appendix C.7 for details and derivations. To draw conclusions about allocative inefficiencies, we compare the allocation and outcomes in the monopsonistic labor market to those that would arise in a competitive (Walrasian) labor market. By a competitive market, we mean that there are no taxes ($\lambda = \tau = 1$) and that all firms act as price takers, as if they face perfectly elastic labor supply curves. To simplify the expressions, we abstract from the transitory term V_{it} . This simplification does not affect the analysis as V_{it} cancels out because it is paid to the worker independent of the choice of firm.

We begin by examining how the firm's problem and solution differ in the two types of labor markets. In the competitive market, firms will hire workers until the marginal revenue product of labor equals the wage rate. Thus, the first order condition of the firm is given by:

$$W_{jt}^{c}(X) = \underbrace{X^{\theta_{j}}(1 - \alpha_{r(j)})A_{jt}(L_{jt}^{c})^{-\alpha_{r(j)}}}_{\text{marginal product of labor: } \mathcal{M}_{it}^{c}(X)}$$

where the superscript c makes clear that we refer to the competitive market case. By comparison, when the firm behaves as a local monopsonist, the wages are marked down relative to the marginal revenue product of labor.

To understand when, and by how much, firms mark down wages, it is useful to rewrite the

structural wage equation (18) as follows:

$$W_{jt}(X) = \left(\underbrace{1 + \frac{\rho_{r(j)}}{\lambda\beta}}_{\text{labor wedge}}\right)^{-1} \cdot \underbrace{X^{\theta_j}(1 - \alpha_{r(j)})A_{jt}L_{jt}^{-\alpha_{r(j)}}}_{\text{marginal product of labor: } \mathcal{M}_{jt}(X)}.$$
(25)

Consider the special case when ρ_r is zero for every r, which means that workers view all firms within each market as perfect substitutes. As a result, firms would face horizontal labor supply curves. Thus, wages would be paid at a level equal to the marginal revenue product of labor, implying no labor wedges. If, on the other hand, ρ_r differs from zero, then firms within the same market are viewed as imperfect substitutes by the workers. Thus, they have wage setting power and workers are paid less than their marginal product. The steeper the labor supply curve facing the firm, the larger the labor wedges. The labor supply curve is steeper, the more important the amenities are for worker's choice of firms. Therefore, labor wedges increase in the variability of the idiosyncratic tastes β^{-1} and in the progressivity of labor income taxes $(1 - \lambda)$.

We now shift attention to the worker's problem and solution in the two types of labor markets. But first, it is useful to express wages using the wage index $I_{rt}(X)$ from equation (16) in the following way:

$$W_{jt}(X) = \underbrace{\frac{W_{jt}(X)}{I_{r(j)t}(X)}}_{\tilde{\mathcal{W}}_{jt}(X)} \underbrace{\frac{I_{r(j)t}(X)}{\sum_{r'} I_{r't}(X)}}_{\bar{\mathcal{W}}_{rt}(X)} \sum_{r'} I_{r't}(X)$$
(26)

Equation (26) allows us to derive parsimonious expressions for the supply of workers to markets and to firms within a market. In the competitive labor market, we get

$$\begin{split} \bar{S}_{rt}^{\mathrm{c}}(X) &= N \Pr[j(i,t) \in J_r | \left\{ W_{jt}^{\mathrm{c}}(X) \right\}_{j=1,\dots,J}, X] \\ &= N \bar{\mathcal{W}}_{rt}^{\mathrm{c}}(X)^{\beta} \\ \tilde{S}_{jt}^{\mathrm{c}}(X) &= \Pr[j(i,t) = j | \left\{ W_{jt}^{\mathrm{c}}(X) \right\}_{j=1,\dots,J}, X, j \in J_r] \\ &= G_j(X)^{\beta/\rho_{r(j)}} \tilde{\mathcal{W}}_{it}^{\mathrm{c}}(X)^{\beta/\rho_{r(j)}} \end{split}$$

and in the monopsonistic labor market we obtain:

$$\bar{S}_{rt}(X) = N \Pr[j(i,t) \in J_r | \{W_{jt}(X)\}_{j=1,\dots,J}, X]$$
$$= N \bar{\mathcal{W}}_{rt}(X)^{\lambda\beta}$$
$$\tilde{S}_{jt}(X) = \Pr[j(i,t) = j | \{W_{jt}(X)\}_{j=1,\dots,J}, X, j \in J_r]$$
$$= G_i(X)^{\beta/\rho_{r(j)}} \tilde{\mathcal{W}}_{it}(X)^{\lambda\beta/\rho_{r(j)}}$$

Comparing these expressions shows the tax wedge in the worker's choice of where to work. Progressive taxation (i.e. $\lambda < 1$) may distort the worker's decision of which market and firm to work in. The reason is that wages are taxed but amenities are not. Thus, progressive taxation makes amenities relatively more important for the workers' choice of where to work.

While the discussion above clarifies the sources of wedges in the problems of the worker and the firm, it does not directly speak to the questions of whether, and in what situations, the equilibrium allocation of workers to firms will be inefficient. To answer these questions, we need to examine how the labor and tax wedges play out in equilibrium. To do so, it is useful to express the marginal product of labor as

$$\mathcal{M}_{jt}(X) = \underbrace{\frac{\mathcal{M}_{jt}(X)}{I_{r(j)t}^{\mathcal{M}}(X)}}_{\tilde{\mathcal{M}}_{jt}(X)} \underbrace{\frac{I_{r(j)t}^{\mathcal{M}}(X)}{\sum_{r'} I_{r't}^{\mathcal{M}}(X)}}_{\bar{\mathcal{M}}_{rt}(X)} \sum_{r'} I_{r't}^{\mathcal{M}}(X)$$
(27)

where

$$I_{rt}^{\mathcal{M}}(X) \equiv \left(\sum_{j' \in J_r} \left(\tau^{1/\lambda} G_{j'}(X)^{1/\lambda} \mathcal{M}_{j't}(X)\right)^{\lambda \beta/\rho_r}\right)^{\rho_r}$$

Next, we equate demand and supply between and within markets. In the competitive labor market, we get

$$\bar{D}_{rt}^c(X) = \bar{S}_{rt}^c(X) = \bar{\mathcal{M}}_{rt}^{c\beta}$$

$$\tilde{D}_{jt}^c(X) = \tilde{S}_{jt}^c(X) = G_j(X)^{\beta/\rho_{r(j)}} \tilde{\mathcal{M}}_{jt}^{c\beta/\rho_{r(j)}}$$
(28)

while in the monopsonistic labor market we obtain:

$$\bar{D}_{rt}(X) = \bar{S}_{rt}(X) = \left(1 + \frac{\rho_r}{\lambda\beta}\right)^{-\lambda\beta} \bar{\mathcal{M}}_{rt}^{\lambda\beta}$$

$$\tilde{D}_{jt}(X) = \tilde{S}_{jt}(X) = G_j(X)^{\beta/\rho_{r(j)}} \tilde{\mathcal{M}}_{jt}^{\lambda\beta/\rho_{r(j)}}$$
(29)

Comparing these expressions allows us to draw inferences about allocative inefficiencies within and between markets.

Within each market, the tax wedge that arise because $\lambda < 1$ is the only source of allocative inefficiency, distorting the worker's ranking of firms in favor of those with better amenities. As a result, these firms can hire workers at relatively low wages, and, therefore, get too many workers as compared to the allocation in the competitive labor market.

Between markets, allocative inefficiencies may arise not only because of the tax wedge but also due to differences in labor wedges across markets. To understand the latter source of inefficiencies, consider the special case when $\lambda = 1, \beta > 0$ and ρ_r is non-zero but the same across all markets. In this case, taxes are proportional but there are still labor wedges and rents in the economy. However, the labor wedges will be the same across all markets. As a consequence, the monopsonistic market allocation of workers to firms is identical to the allocation one would obtain in the competitive equilibrium. A corollary of this results is that tax wedges are the only source of allocative inefficiencies if one assumes a standard logit structure on the distribution of $\vec{\epsilon}_{it}$ (as in, for example, Card et al., 2018).

With the nested logit structure on the distribution of $\vec{\epsilon}_{it}$, allocative inefficiencies across markets may arise because ρ_r can vary across markets, implying that workers may view firms as closer substitutes in some markets than others. This will create differences across markets in the wage setting power of firms, and so in their abilities to mark down wages. Markets facing an elastic labor supply curve (i.e. low value of ρ_r) will have relatively high wages and, as a result, attract too many workers compared to the allocation in the competitive equilibrium. Progressive taxation will amplify any differences in ρ_r across markets, leading to an even larger misallocation of workers to firms.

To improve the allocation of workers to firms, the government can change the tax system in two ways. First, a less progressive tax system (i.e. increase λ) may reduce the misallocation that arises from the tax wedge. Second, letting τ vary across markets may improve the allocation of workers by counteracting differences in the wage setting power of firms. After estimating the parameters of the model, we perform counterfactuals that quantifies the impacts of such tax reforms on the equilibrium allocation and outcomes, including earnings, output and welfare.

5 Identification and estimation

We now describe how the quantities of interest are identified and estimated. We provide a formal identification argument while summarizing, in Table 6, the parameters needed to recover a given quantity of interest and the moments used to identify these parameters. Our identification argument reveals that many of these quantities do not require knowledge of all the structural parameters. Thus, some of our findings may be considered more reliable than others.

5.1 Parameters needed to recover rents of workers and employers

As shown in Results 1, 2 and 4, we only need to know $(\alpha_r, \beta, \rho_r)$ to draw inference about the expected rents of workers and employers at the firm and market level. Our identification argument therefore proceeds by showing how these parameters can be identified from the panel data on workers and firms. However, before we present the formal identification argument, it is useful to consider what one can and cannot identify directly from an ideal experiment. This consideration clarifies the necessary assumptions even with an ideal experiment and the additional ones needed because we do not have such an experiment.

Ideal experiment

Suppose we observed and were able to exogenously change both \tilde{a} , the component of productivity that is specific to a firm, and \bar{a} , the component of productivity that is common to all firms in a market. As evident from the structural equations (18)-(20), exogenous changes in \tilde{a} and \bar{a} affect both the firm's value added $y_i(\bar{a}, \tilde{a})$ and the wages it offers to workers of a given quality $w_j(x, v, \bar{a}, \tilde{a})$:

$$\frac{\partial w_j(x,v,\bar{a},\tilde{a})}{\partial \tilde{a}} \left(\frac{\partial y_j(\bar{a},\tilde{a})}{\partial \tilde{a}}\right)^{-1} = \frac{1}{1+\lambda\beta/\rho_{r(j)}}$$
$$\frac{\partial w_j(x,v,\bar{a},\tilde{a})}{\partial \bar{a}} \left(\frac{\partial y_j(\bar{a},\tilde{a})}{\partial \bar{a}}\right)^{-1} = \frac{1}{1+\lambda\beta}$$

Since λ is a known (or pre-estimated) tax parameter, β and $\rho_{r(j)}$ can be identified from these two equations.

In this ideal experiment, the pass-through to $w_j(x, v, \bar{a}, \tilde{a})$ of an \bar{a} induced change in the $y_j(\bar{a}, \tilde{a})$ would identify β . Given this parameter, the pass-through to $w_j(x, v, \bar{a}, \tilde{a})$ of an \tilde{a} induced change in value added would identify $\rho_{r(j)}$. Next, one could recover α_r by combining the structural equations (18)-(20) to get:

$$\mathbb{E}\left[y_{jt} - b_{jt} | j \in J_r\right] = -c_r = -\log(1 - \alpha_r) - \log\left(\frac{\lambda\beta/\rho_r}{1 + \lambda\beta/\rho_r}\right).$$
(30)

This equation shows that the market-specific scale parameter can be recovered from data on wage bills and value added per market, after correcting for labor wedges as captured by the term $\log\left(\frac{\lambda\beta/\rho_r}{1+\lambda\beta/\rho_r}\right)$.

It is important to observe that performing such an ideal experiment is difficult for several reasons. First, it can be difficult to find shifters or instruments for \tilde{a} and \bar{a} that are both as good as randomly assigned to firms and have skill-neutral impacts on productivity. Existing studies using instruments have either a shifter for \tilde{a} or a shifter for \bar{a} , but not both.¹¹ Second, even if the analyst finds such instruments, it is necessary to keep amenities fixed. In other words, firms must not change the amenities offered to workers in response to shifts in productivity. Otherwise, value added would not be the only channel through which \tilde{a} and \bar{a} would affect wages. Third, the analyst needs to observe worker quality or make restrictions on worker mobility in response to shocks to firm productivity. Otherwise, the productivity effects on the wages the firm pays to workers of a given quality may be confounded by shifts in the quality of the workforce of the firm.

Panel data approach

Since we do not have such an ideal experiment, we instead have to rely on panel data on workers and firms. We now show how these data can be used to identify the structural parameters $(\alpha_r, \beta, \rho_r)$. We begin by describing the restrictions on the primitives that deliver the statistical processes for wages and value added. This forges a direct link between our model of the labor market in Section 4 and the pass through analysis in Section 3.2. Once this link is established,

¹¹See, for example, Abowd and Lemieux (1993), Van Reenen (1996), Kline et al. (2018a), Garin and Silvério (2019), and Berger et al. (2019). If one only has an instrument for \tilde{a} , it would still be possible to achieve identification by assuming that ρ_r is zero for every market r, which means that workers view all firms within each market as perfect substitutes. However, our data are at odds with this assumption. If, on the other hand, one only has an instrument for \bar{a} , then identifying the labor supply curve facing the firm seems difficult, even with additional assumptions.

we proceed by showing how the estimates from the pass-through analysis can be used to recover the structural parameters.

The starting point is the structural equations (18)-(20) that express wages $w_j(x, v, \bar{a}, \tilde{a})$ and value added $y_j(\bar{a}, \tilde{a})$ as functions of primitives that are fixed over time $\Gamma = (\bar{p}_r, \tilde{p}_j, g_j(x), x_i)$ and those that vary over time $(\bar{z}_{rt}, \tilde{z}_{jt}, v_{it})$. We begin by specifying how the time-varying components evolve. The productivity of the worker is assumed to evolve according to the following process:

$$v_{it} = v_{it}^{\mathrm{p}} + \xi_{it}^{x} + \delta^{x}\xi_{it-1}^{x}$$
$$v_{it}^{\mathrm{p}} = v_{it-1}^{\mathrm{p}} + \mu_{it}$$

where the persistent component v_{it}^{p} is a unit root process with innovation shocks μ_{it} , and the transitory component is assumed to follow a MA(1) process with coefficients δ^{x} and innovation variance σ_{x}^{2} . We assume that firm productivity evolves as a unit root process at both the firm level:

$$\begin{aligned} \tilde{a}_{jt} &= \tilde{p}_j + \tilde{z}_{jt} \\ \tilde{z}_{jt} &= \tilde{z}_{jt-1} + \tilde{u}_{jt} \end{aligned}$$

and the market level:

$$\bar{a}_{rt} = \bar{p}_r + \bar{z}_{rt}$$

$$\bar{z}_{rt} = \bar{z}_{rt-1} + \bar{u}_{rt}$$

On top of this, we allow for measurement error in the observed value added in the form of a transitory component that follows an MA(1) process. This gives the following relationship between value added in the data y_{jt} and value added in the model y_{jt}^* :

$$y_{jt} = y_{jt}^* + \xi_{jt}^y + \delta^y \xi_{jt-1}^y$$

where δ^y is the MA coefficient and σ_y^2 is variance of ξ_{jt}^y .

After specifying these processes, it is necessary to invoke some restrictions on the relationship between the primitives. Letting the history of time-varying unobservables at time t be denoted by $\Omega_t \equiv \{\tilde{u}_{jt'}, \bar{u}_{rt'}, \xi_{it'}^x, \xi_{jt'}^y, \mu_{it'}, \epsilon_{ijt'}\}_{i,j,r,t' \leq t}$, we make the following assumption:

Assumption 3. Assume that each of the random variables $\tilde{u}_{jt'}, \bar{u}_{rt'}, \xi^x_{it'}, \xi^y_{jt'}, \mu_{it}$ are drawn *i.i.d.* conditional on Ω_{t-1} and Γ .

It is important to observe what is and is not being restricted by Assumption 3. It permits arbitrary correlation between the components of Γ . Thus, our model allows for rich heterogeneity of both firms and workers, and systematic sorting of different workers into different firms. However, Assumption 3 implies that worker-specific innovations to productivity are both independent across coworkers and orthogonal to innovations to firm productivity, as well as orthogonal to idiosyncratic taste realizations ϵ_{ijt} . Moreover, worker-specific innovations to productivity do not induce mobility because they are paid to the worker independent of the choice of firm. This is key to identifying the pass-through rates of firm shocks by looking at changes over time in the earnings of incumbent workers. Another important implication of Assumption 3 is that innovations to firm productivity are independent of the endowment of firm amenities. This is crucial for learning about innovations to the firm's productivity from changes in its value added over time.

In Online Appendix D.1, we show that Assumption 3 together with the structural equations (18)-(20) imply Assumptions 1 and 2 and deliver the following relationship between (γ, Υ) and (β, ρ_r) :

$$\gamma = \frac{1}{1 + \lambda \beta / \rho_r}$$
$$\Upsilon = \frac{1}{1 + \lambda \beta}$$

This establishes that the parameters (β, ρ_r) are identified under Assumption 3. Furthermore, it shows how the model of the labor market provides a structural interpretation of the passthrough rates reported in Section 3.2 in terms of expected worker rents at the firm level and at the market level. Once (β, ρ_r) are known, we can recover α_r (and therefore, the expected firm rents) from data on wage bills and value added per market, as shown in equation (30)

5.2 Parameters needed to economically interpret AKM

The empirical analysis in Section 3.3 decomposes earnings inequality into components that are meant to capture worker heterogeneity, firm wage premiums and worker sorting. Our model offers an economic interpretation of this statistical decomposition. For this purpose, we not only need to know $(\alpha_r, \beta, \rho_r)$ but also $(\theta_j, \bar{p}_r, \tilde{p}_j, \sigma_{\tilde{u}}^2, \sigma_{\tilde{u}}^2, h_j, x_i)$.

Given the above restrictions on the primitives, we can recover θ_j and a firm-specific time invariant intercept ψ_j from the analysis of cross-sectional earnings inequality in Section 3.3. As in this analysis, we remove the time-varying firm- and market-specific components of earnings, so that one can express the expected earnings of worker *i* in firm *j* as a the sum of her own productivity x_i times the firm-specific technology parameter θ_j and a firm-specific component ψ_j :

$$\mathbb{E}\left[\left.w_{it} - \frac{1}{1+\lambda\beta}\bar{y}_{rt} - \frac{\rho_r}{\rho_r + \lambda\beta}(y_{jt} - \bar{y}_{rt})\right| j(i,t) = j, \ j \in J_r\right] = \theta_j x_i + \psi_j \tag{31}$$

where for $j \in J_r$ we define:

$$\psi_j \equiv c_r - \alpha_r h_j - \frac{\lambda\beta(\rho_r - 1)(1 - \alpha_r)}{(1 + \lambda\beta)(\rho_r + \beta)}\bar{h}_r.$$
(32)

This derivation is presented in Online Appendix D.2. The firm-specific component ψ_j depends only on the amenities components, h_j and $\bar{h}_r = \mathbb{E}[h_j | j \in J_r]$. In contrast to ψ_{jt} defined before, this ψ_j does not depend on time t firm level and market level productivities \bar{a}_{rt} and \tilde{a}_{jt} . Under the rank condition that workers moving to a firm are not of the exact same quality as workers moving from the firm:

$$\mathbb{E}[x_i|j(i,t)=j, j(i,t+1)=j'] \neq \mathbb{E}[x_i|j(i,t)=j', j(i,t+1)=j]$$

we can identify both θ_j and ψ_j from data on the changes in earnings associated with these moves. Given these parameters, we can recover worker quality x_i from equation (31) and h_j from equation (32).

Knowing x_i , h_j and θ_j for each worker and firm, we can recover the productivity components \tilde{a}_{jt} and \bar{a}_{rt} from equation (21) which defines the total efficiency unit of labor:

$$\ell_{jt} = \log \sum_{i \text{ st } j(i,t)=j} X_i^{\theta_j}$$

= $h_j + \frac{\lambda \beta / \rho_{r(j)}}{1 + \alpha_{r(j)} \lambda \beta / \rho_{r(j)}} \tilde{a}_{jt} + \frac{\lambda \beta}{1 + \alpha_{r(j)} \lambda \beta} \bar{a}_{r(j)t}$

The intuition is straightforward. Conditional on the amenity component h_j , we recover firm productivity \tilde{a}_{jt} by comparing ℓ_{jt} across firms within the same market. Holding market amenities fixed, we identify \bar{a}_{rt} by comparing total efficiency units of labor of different markets. Once we know $(\bar{a}_{rt}, \tilde{a}_{jt})$, we can recover $(\bar{p}_r, \tilde{p}_j, \sigma_{\tilde{u}}^2, \sigma_{\tilde{u}}^2)$ from the assumptions about how productivity evolves at the firm and market level.

5.3 Parameters needed for welfare and counterfactuals

To make inference about welfare and to perform counterfactuals, it is necessary to also recover the preferences component $G_j(X)$. This is done through a revealed preference argument: Holding wages fixed, firms with favorable amenities (for a given type of worker) are able to attract more workers (of that type). Conditional on wages, the size and composition of firms and markets should therefore be informative about unobserved amenities.

We formalize this intuition in Lemma 13 in Online Appendix D.3, showing that $G_j(X)$ can be identified from data on the allocation of workers to firms and markets. We first derive the following expression which links $G_j(X)$ to the size and composition of firms and markets:

$$\exp(\lambda\psi_{jt})X^{\lambda\theta_j}G_j(X) = \left(\Pr\left[j(i,t)\in J_{r(j)}\mid X\right]\right)^{1/\beta} \left(\Pr\left[j(i,t)=j\mid X, r=r(j)\right]\right)^{\rho_{r(j)}/\beta} \qquad \forall j,t,X$$

where $Pr\left[j(i,t)\in J_{r(j)} \mid X\right]$ is the probability that workers of quality X choose to work in market r, while $Pr\left[j(i,t)=j \mid X, r=r(j)\right]$ is the probability these workers choose to work for firm j conditional on selecting market r. Consider two firms j and j' in the same market r. The differences in size and composition of these firms depend on the gaps in wages and amenities:

$$\underbrace{\lambda\left((\theta_{j} - \theta_{j'})x_{i} + \psi_{j} - \psi_{j'}\right)}_{\text{wage gap}} + \underbrace{\log G_{j}(X) - \log G_{j'}(X)}_{\text{amenity gap}} = \underbrace{\frac{\rho_{r}}{\beta}}_{\text{within-market elasticity}} \underbrace{\log \frac{\Pr\left[j(i, t) = j \mid X, r(j) = r\right]}{\Pr\left[j(i, t) = j' \mid X, r(j') = r\right]}}_{\text{relative size by worker type}}$$

Since both the wage gap and the within-market elasticity are already identified, we can recover the value of amenities up to a common market factor by comparing the size and composition of firms. Comparing firms across markets allow us to pin down the market factor.

5.4 Estimation

The estimation procedure follows closely the identification arguments laid out above. In the estimation, however, we impose a few additional restrictions on the heterogeneity of workers, firms and markets. Neither of these restrictions are necessary for identification, but they help reduce the number of parameters to estimate. We now describe these restrictions before presenting the parameter estimates and assessing the fit of the model.

Specification of empirical model

We begin by restricting the market-specific parameters α_r and ρ_r to be the same within broad markets (as defined in Section 2). The restriction on α_r means the scale parameter can vary freely across, but not within, broad regions and sectors of the economy. The assumption on ρ_r restricts the nested logit structure of the preferences. Recall that the parameter ρ_r measures the degree of independence in a worker's taste for alternative firms within the nest. We specified the nest as the combination of commuting zone and two-digit industry. We now restrict the parameter ρ_r to be the same for all nests within each broad market. As a result, labor wedges may vary across, but not within, broad regions and sectors.

A second set of restrictions is that we draw the firm-specific components θ_j and ψ_j from a discrete distribution. We follow Bonhomme et al. (2019) in using a two-step grouped fixedeffects estimation, which consists of a classification and an estimation step. In a first step, firms are classified into groups indexed by k based on the empirical earnings distribution using the k-means clustering algorithm. The k-means classification groups together firms whose earnings distributions are most similar.¹² Then, in a second step, we estimate the parameters $\theta_{k(j)}$ and $\psi_{k(j)}$. In the baseline specification, we assume there exist 10 firm types. We view the assumption of discrete heterogeneity as a technique for dimensionality reduction in the estimation, instead of viewing discreteness as a substantive assumption about population unobservables. It is reassuring to find that the estimates of firm effects do not change materially if we instead allow for

 $^{^{12}}$ Here, we follow Bonhomme et al. (2019). Concretely, we use a weighted k-means algorithm with 100 randomly generated starting values. We use the firms' empirical distributions of log-earnings on a grid of 10 percentiles of the overall log-earnings distribution.

20, 30, 40 or 50 firm types (see e.g. Online Appendix Figure B.4).

Lastly, we also make the following discreteness assumption for the systematic components of firm amenities:

$$G_j(X) = \bar{G}_{r(j)}\tilde{G}_j G_{k(j)}(X),$$

where we define the firm class k(j) within market r using the classification of Section 3 interacted with the market. This multiplicative structure reduces the number of parameters we need to estimate while allowing for systematic differences in amenities across firms and markets $(\tilde{G}_j, \bar{G}_{r(j)})$ and heterogeneous tastes according to the quality of the worker $G_{k(j)}(X)$. As a result, amenities may still generate sorting of better workers to productive firms, and compensating differentials may still vary across firms, markets and workers, and heterogeneity in amenities. For estimation purposes, we take advantage of the derivations in Online Appendix D.3, which express the preference components $(\bar{G}_{r(j)}\tilde{G}_jG_{k(j)}(X))$ as functions of the size and composition of firms and markets:

$$G_{k}(X) = X^{-\lambda\theta_{k}} \left(\frac{Pr[X|r]}{Pr[X]}\right)^{1/\beta} \left(\frac{Pr\left[X|k\right]}{Pr\left[X|r\right]}\right)^{\rho_{r}/\beta}$$

$$Pr[j(i,t)=j \mid j(i,t)\in J_{r}] = \tilde{G}_{j}^{\beta/\rho_{r}} \int \frac{\left(\tau G_{k(j)}(X)\exp(\lambda\psi_{jt})X^{\lambda\theta_{j}}\right)^{\beta/\rho_{r}}}{\sum_{j'\in J_{r}} \left(\tau \tilde{G}_{j'}G_{k(j')}(X)\exp(\lambda\psi_{j't})X^{\lambda\theta_{j'}}\right)^{\beta/\rho_{r}}} Pr[X|r]dX$$

$$Pr[j(i,t)\in J_{r}\mid X] = \bar{G}_{r}^{\beta} \int \frac{\left(\sum_{j'\in J_{r}} \left(\tau \tilde{G}_{j}G_{k(j)}(X)\exp(\lambda\psi_{j't})X^{\lambda\theta_{j'}}\right)^{\beta/\rho_{r}}\right)^{\rho_{r}} NM_{X}(X)}{\sum_{r'} \left(\sum_{j'\in J_{r'}} \left(\tau \bar{G}_{r'}\tilde{G}_{j'}G_{k(j')}(X)\exp(\lambda\psi_{j't})X^{\lambda\theta_{j'}}\right)^{\beta/\rho_{r'}}\right)^{\rho_{r'}} dX$$

where $\tilde{G}_{j}^{\beta/\rho}$ and $\bar{G}_{r(j)}^{\beta}$ are estimated by iterating on the expression above. In this estimation of $G_{j}(X)$, we rely on a discretization of X using 10 points of support and we also group markets into 10 different market types based on their realized empirical distribution of earnings.

Estimates and fit

Consider first the parameters of the tax function, τ and λ . We estimate these parameters outside the model. In each year, we regress log net household income (earnings plus other income minus taxes) on log household gross income (earnings plus other income) for our sample. The construction of these income measures is detailed in Online Appendix A. The intercept from this regression gives us τ while λ is identified from the slope coefficient. We estimate τ of around 0.89 whereas λ is estimated to be about 0.92.¹³ In a proportional tax-transfer system, λ is equal to one and $(1 - \tau)$ is the proportional effective tax rate. By contrast, if $0 < \lambda < 1$, then the marginal effective tax rate is increasing in earnings. Online Appendix Figure B.6 shows how well our parsimonious tax function approximates the effective tax rates implicit in the complex U.S. tax-transfer system. Here we compare the predicted log net income from the regression

¹³These results mirror closely existing U.S. estimates of τ and λ (see e.g. Guner et al., 2014, Heathcote et al., 2017).

to the observed log net income across the distribution of log gross income, finding that this specification provides an excellent fit.

Once we have the tax parameters, we can estimate the other parameters following the identification arguments above. In other words, we substitute the sample moments in place of the population moments. Table 6 summarizes the estimates of key parameters. On average, we find α_r to be around 0.21, suggesting declining returns to labor input. The estimate of β is about 4.99. This finding suggests considerable variability across workers in the idiosyncratic taste for a given firm. On average, we estimate ρ_r to be about 0.70. This implies a correlation of 0.51 in the taste of the worker across firms in the same industry and location. The variation in worker quality is empirically important, whereas the time-varying firm premia vary less. Firms differ significantly in their technology, productivity and amenities. A considerable part of the productivity dispersion is across markets.

Because we allow for a lot of heterogeneity across workers and firms, our model fits perfectly many aspects of the data. However, some moments of the data are not necessarily matched that well. In Figure 4, we compare the observed and the predicted values of firm effects, value added, efficiency units of labor, and wage bill. We make this comparison separately according to the actual and predicted firm size. When considering these moments (which are not directly targeted), the model performs well.

In Figure 5, we take advantage of the fact that the amenity component h_j is over-identified. One possibility is to recover it from the equation for firm wage premiums (32), as we did in the estimation. Another possibility is to use the fixed-point definition of h_j as a function of $(\tilde{P}_j, \bar{P}_r, G_j(X))$, as shown in Lemma 8. This definition comes from the equilibrium constraint of the model, which we do not directly use in the estimation. Figure 5 shows that the values of h_j based on equation (32) are very similar to those we obtain from the equilibrium constraint of the model. This finding increases our confidence in the moment conditions implied by our economic model.

6 Empirical insights from the model

We now present the empirical insights from the estimated model. These insights require an explicit model of the labor market, and, thus, they may be susceptible to model misspecification. As shown in Section 5, however, many of the insights do not require knowledge of all the structural parameters. Thus, some of our findings may be considered more reliable than others. To make this clear, we first present the findings that rely on the least assumptions and then move to those that require additional restrictions on the functioning of the labor market.

6.1 Rents and and labor wedges

Table 7 presents estimates of the size of rents earned by American firms and workers from ongoing employment relationships. In the first column of the table, we estimate these quantities under the restrictions that α_r is the same for each broad market and that ρ_r is equal to 1 for

Name		Unique Parameters	Mean Estimate	Moments	of the Data
Panel A. Rents and Scale					
Idiosyncratic Taste Parameter	β	1	4.99	Market Passthrough	$\frac{\mathbb{E}\left[\Delta \bar{y}_{rt}(\bar{w}_{rt+\tau} - \bar{w}_{rt-\tau'}) S_i=1\right]}{\mathbb{E}\left[\Delta \bar{y}_{rt}(\bar{y}_{rt+\tau} - \bar{y}_{rt-\tau'}) S_i=1\right]}$
Taste Correlation Parameter	$ ho_r$	×	0.70	Net Passthrough	$\mathbb{E}\left[\Delta \tilde{y}_{j,t}\left(\tilde{w}_{it+\tau} - \tilde{w}_{it-\tau'}\right) S_i=1, r(j)=r\right]$ $\mathbb{E}\left[\Delta \tilde{y}_{i,t}\left(\tilde{y}_{i,t+\tau} - \tilde{y}_{i,t-\tau'}\right) S_i=1, r(i)=r\right]$
Returns to Scale Parameter	α_r	×	0.21	Labor Share	$\mathbb{E}[b_{j(i,t)} - y_{j(i,t)} r(j) = r]$
Name		Unique Parameters	Var. Estimate	Moments	of the Data
Panel B. Economic Interpretation of	AKM				
Time-varying Firm Premium	ψ_{jt}	10,669,602	0.02	Structural Wage Equation	$\mathbb{E}[w_{it} - rac{1}{1+\lambdaeta}ar{y}_{r,t}]$
Firm-specific Technology Parameter	$ heta_{j}$	10	0.04		$-rac{ ho_{r,+}}{ ho_{r,+}+\lambdaeta} ilde{y}_{j,t} r(j)=r]$
Worker Quality	x_i	61,670,459	0.31	Wage Changes around Moves	$\mathbb{E}[w_{it+1} j \rightarrow j'] - \mathbb{E}[w_{it} j' \rightarrow j]$
Amenity Efficiency Units at Neutral TFP	h_{j}	1,953,915	0.14		$\mathbb{E}[w_{it} j' \to j] - \mathbb{E}[w_{it+1} j \to j']$
Rirm-snecific TFP	ñ	1 953 915	0.04	Total Labor Innut ℓ_r	$I_{i,i} \equiv \log \sum X^{ heta_j}$ and $\psi_{i,i}$
Market-specific TFP	\bar{p}_r	114,773	0.12	Time-varying Firm Premium	if a new interpose of
Name		Unique Parameters	Var. Estimate	Moments	of the Data
Panel C. Model Counterfactuals Preferences for amenities for:	$g_j(X)$	37,236,342	0.20	Firm Size &	$\Pr[j]$
Firm j for workers of quality X Market r for workers of quality X				Firm Composition $\&$ Market Composition	$\Pr[x k(j) = k]$ $\Pr[x r(j) = r]$

l Moments
Targetec
Parameters and
Model
f Interest,
Quantities o
Table 6:

Notes: This table displays the model parameters and the moments targeted in their estimation.



Figure 4: Fit of the Model for Untargeted Moments

Notes: In this figure, we compare the observed and the predicted values of firm effects, value added, efficiency units of labor, and wage bill. We make this comparison separately according to the actual and predicted firm size.

all firms in every market. In the second and third column we relax these restrictions, allowing for the possibility that technology differs across markets and that workers view firms within the same market as closer substitutes than firms in different markets. To easily compare the estimates from the baseline specification (in the second and third column) to those produced by the restricted specification (in the first column), we report national averages and refer to Online Appendix Table B.3 for the market-specific results.

In both specifications, we find evidence of significant amount of rents and imperfect competition in the U.S. labor market due to the horizontal employer differentiation. When we impose the restrictions on α_r and ρ_r , we estimate that workers are, on average, willing to pay 14 percent of their annual earnings to stay in their current jobs. This corresponds to \$5,875 per worker. By comparison, firms earn, on average, 11 percent of profits from rents (with profits being measured as value added minus the wage bill). This amounts to \$5,932 per worker in the firm. These results imply that the rents from imperfect competition in the labor market are split equally between the employers and the workers.

Comparing the results across the columns of Table 7 shows that relaxing the restrictions on ρ_r and α_r does not materially change the estimates of firm level rents or the workers' share of



Figure 5: Estimates of the Amenity Components h_j from the Wage Equation versus the Equilibrium Constraint

Notes: In this figure, we plot the mean of h_j across log size bins. We compare the baseline estimates of h_j from the equation for firm wage premiums (32), versus those estimated using the equilibrium constraint by solving the fixed-point definition of h_j as a function of $(\tilde{P}_j, \bar{P}_r, G_j(X))$, as shown in Lemma 8.

total rents. Yet relaxing these restrictions is still important for the conclusions we draw about the labor market. One insight from the second and third column is that the market level rents are considerably larger than the firm level rents. The workers are, on average, willing to pay \$7,331 (18 percent of their annual earnings) to avoid having to work for a firm in a different market. By way of comparison, the average willingness to pay to remain in the current firm is only \$5,447 (13 percent of earnings). The relatively large market level rents reflect that firms within the same market are more likely to be close substitutes than firms in different markets.

By relaxing the restrictions on ρ_r , we also learn that labor wedges are significant and vary substantially across markets. In Figure 6, we report the estimates of labor wedges from the baseline specification. On average, the marginal revenue product of labor is 15 percent higher than the wage. Behind this average, however, there is important variation. Empirically, the labor wedges are most pronounced in the goods sector (which have higher values of ρ_r). In the West, for example, the labor wedge is 6 percentage points larger for firms in the goods sector as compared to those in the service sectors.

6.2 Compensating differentials

The estimates of rents suggest the average American worker is far from the margin of indifference in his choice of firm, and would, in fact, maintain the same choice even if his current firm offered significantly lower wages. In other words, the average worker would, if asked, identify amenities as important to his choice of firm. This finding does not, however, imply that the marginal

	Ren	ts and Rent-	shares
	Firm Only Firm-level	Accounting Firm-level	g for Markets Market-level
Workers' Rents:			
Per-worker Dollars	5,875 (284)	5,447 (395)	7,331 (1,234)
Share of Earnings	14% (1%)	13% (1%)	18% (3%)
Firms' Rents:			
Per-worker Dollars	5,932 (709)	$5,780 \\ (1,547)$	7,910 (1,737)
Share of Profits	11% (1%)	11% (3%)	15% (3%)
Workers' Share of Rents	50% (2%)	49% (4%)	48% (3%)

Table 7: Estimates of rents and rent sharing (national averages)

Notes: This table displays our main results on rents and rent-sharing. Column 1 presents results from the specification which imposes $\Upsilon = \gamma$, $\rho_r = 1$, and $\alpha_r = \alpha$ ("Firm only"), while columns 2-3 report results from the specification which allows Υ to differ from γ , and for ρ_r and α_r to vary across broad markets ("Accounting for Markets"). Standard errors are estimated using 40 block bootstrap draws in which the block is taken to be the market.

workers view the amenities of the current firm as much better or much worse than those offered by other firms. To gauge the preferences for amenities of the marginal workers, we estimate the average or expected compensated differential.

The estimates of the expected compensating differentials are displayed in Figure 7. To estimate these quantities, we randomly draw two firms, j and j', from the overall distribution of firms (where each firm is drawn with probability proportional to its size). We then compute the compensating differentials between j and j' for a worker of given quality x as $\psi_j + x\theta_j - \psi_{j'} - x\theta_{j'}$. We repeat this procedure for a large number of draws of firms.

The solid horizontal line in Figure 7 shows the average or expected compensating differential for a marginal worker. For two randomly drawn firms, the one with worse amenities can be expected to pay an additional 18 percent in order to convince marginal workers (of average quality) to accept the job. There is, however, considerable heterogeneity in the compensating differentials according to worker quality. The upward sloping solid line shows how the expected compensating differential varies with worker quality. For high quality workers (95 percentile in the national distribution), the expected compensating differentials are as large as 30 percent. By comparison, marginal workers of low quality (5 percentile in the national distribution) require less than 10 percent additional pay to work in the firm with the unfavorable amenities.

The dashed lines of Figure 7 display the compensating differentials across firms within a



Figure 6: Labor Wedges

Notes: In this figure, we display the estimated labor wedge for each of the 8 broad markets. The population-weighted mean labor wedge is represented by a horizontal line.

market. To compute these quantities, we use the same procedure as above, except we now compare firms within each market. For two randomly drawn firms in the same market, the one with worse amenities can be expected to pay an additional 14 percent in order to convince marginal workers (of average quality) to accept the job. This suggests that three-quarters of compensating differentials reflect differences in amenities within markets.

6.3 Understanding the firm effects

Why is the inequality contribution from firm effects so small? To answer this question, it is necessary to understand the economic determinants of the firm component of the AKM model. As evident from equation (18) of Section 4.3, our structural wage equation reduces to the AKM model if there are no production complementarities and every firm faces a perfectly elastic labor supply curve. The structural wage equation would then be log additive in worker and firm effects, with the size of the firm effects reflecting the heterogeneity in amenities across firms.

We find, however, evidence of upward sloping supply curves for labor, which generate a positive relationship between the firm's productivity and the wages it pays. As a result, the size of the firm effects depends not only on the heterogeneity in firm amenities, but also on the differences in productivity across firms as well as the covariance between productivity and amenities within firms:



Figure 7: Compensating differentials

Notes: In this figure, we plot mean compensating differentials overall and within market. To do so, we randomly draw a pair of firms (j, j') with probability proportional to size. Each j' is drawn from the full set of firms when estimating overall compensating differentials and from the set of firms in the same market as j when estimating within market compensating differentials. Then, we estimate the compensating differential between j and j' for a worker of given quality $x_i = x$ by $\psi_j + x\theta_j - \psi_{j'} - x\theta_{j'}$. This figure plots the mean absolute value of the compensating differentials across deciles of the x_i distribution, where the horizontal lines denote means across the distribution of x_i .

$$\begin{aligned} Var(\psi_{j(i,t),t}) &= Var(c_r - \alpha_r h_{j(i,t)} + \frac{1}{1 + \alpha_r \lambda \beta} \bar{a}_{rt} + \frac{1}{1 + \alpha_r \lambda \beta / \rho_r} \tilde{a}_{j(i,t),t}) \\ &= \underbrace{Var(c_r - \alpha_r h_{j(i,t)})}_{\text{Amenities}} + \underbrace{Var(\frac{1}{1 + \alpha_r \lambda \beta} \bar{a}_{rt} + \frac{1}{1 + \alpha_r \lambda \beta / \rho_r} \tilde{a}_{j(i,t),t})}_{\text{TFP}} \\ &+ \underbrace{2Cov(c_r - \alpha_r h_{j(i,t)}, \frac{1}{1 + \alpha_r \lambda \beta} \bar{a}_{rt} + \frac{1}{1 + \alpha_r \lambda \beta / \rho_r} \tilde{a}_{j(i,t),t})}_{\text{covariance between amenities and TFP}} \end{aligned}$$

These components can be broken down between and within broad markets (with broad market being defined as the combination of sector times region). Within these broad markets, the components can be decomposed further by looking within and between detailed markets (where detailed market is defined as the combination of industry times commuting zone).

The results from these decompositions are reported in Table 8. They suggest a lot of variation in amenities and productivity across firms. Interpreted in isolation, this heterogeneity predicts a large inequality contribution from firm effects. However, productive firms tend to have good amenities, which act as compensating differentials and push wages down in productive firms.

	Between Broad Markets	Within Bro	ad Markets
		Between Detailed Markets	Within Detailed Markets
Total	0.4%	2.0%	3.1%
Decomposition:			
Amenity Differences	15.9%	7.8%	7.1%
TFP Differences	15.5%	11.9%	8.6%
Amenity-TFP Covariance	-31.1%	-17.7%	-12.6%

Table 8: Decomposition of the Variation in Firm Premiums

Notes: This table displays our estimates of the decomposition of time-varying firm premium variation in three levels: variation between broad markets, between detailed markets (within broad markets), and between firms (within detailed markets). Broad markets are defined as the combination of sector times region, and detailed markets are defined as the combination of industry times commuting zone. We decompose the variation in time-varying firm premiums into the contributions from amenity differences, TFP differences, and the covariance between amenity and TFP differences. All variances are expressed as shares of log earnings variance.

As a result, firm effects explain only a few percent of the overall variation in log earnings. For example, firm effects within detailed markets explain 3.1 percent of the variation in log earnings, which is much less than predicted by the variances of firm productivity (8.6 percent) and amenities (7.1 percent).

6.4 Determinants of worker sorting

There are several possible reasons why better workers are overrepresented in higher paying firms. One possible reason is that productive firms have better amenities, and high wage workers may value amenities more than low wage workers. Another possible reason is complementarities in production, which incentivizes better workers to sort into productive firms. We now perform counterfactuals that help quantify the importance of these distinct reasons for sorting in terms of the correlation between worker and firm effects (i.e. $Corr(x_i, \psi_{j(i,t)})$) and the inequality contribution from worker sorting (i.e. $2Cov(x_i, \psi_{j(i,t)})$).

In the counterfactuals we consider, we first reduce the heterogeneity across firms in amenities or production complementarities by replacing either $g_j(x)$ with $(1-s) g_j(x) + s\bar{g}_j$ or θ_j with $(1-s) \theta_j + s\bar{\theta}$, where $\bar{g}_j = \mathbb{E}_x [g_j(x)]$, $\bar{\theta} = \mathbb{E} [\theta_j]$. Here, $s \in [0, 1]$ is the shrink rate with s = 0corresponding to the baseline model. By reducing the heterogeneity in production complementarities, we are effectively making amenities more important for the allocation of workers to firm (and vice versa). Once we have chosen the counterfactual values of $g_j(x)$ or θ_j , we solve for the allocation of workers to firms. Keeping the wages fixed at baseline values (s = 0), we then compute $Corr(x_i, \psi_{j(i,t)})$ and $2Cov(x_i, \psi_{j(i,t)})$ at the counterfactual allocation of workers to firms.

Figure 8 illustrates the importance of amenities versus production complementarities for the sorting of workers to firms. In this figure, we report the quality of the workforce by firm type in the baseline economy with s = 0 (subfigure a) and the counterfactual economies with $s = \frac{1}{2}$ for either amenities (subfigure b) or production complementarities (subfigure c). The results from



Figure 8: Actual and counterfactual composition of the workforce by firm types

Notes: In this figure, we reduce the heterogeneity across firms in amenities or production complementarities by replacing either $g_j(x)$ with $(1-s)g_j(x) + s\bar{g}_j$ or θ_j with $(1-s)\theta_j + s\bar{\theta}$, where $\bar{g}_j = \mathbb{E}_x [g_j(x)], \bar{\theta} = \mathbb{E}[\theta_j]$. Here, $s \in [0, 1]$ is the shrink rate with s = 0 corresponding to the baseline model. We report the quality of the workforce by firm type in the baseline economy with s = 0 (subfigure a) and the counterfactual economies with $s = \frac{1}{2}$ for either amenities (subfigure b) or production complementarities (subfigure c).

these counterfactuals are presented in Figure 8 and suggest that production complementarities are the key reason why better workers are sorting into higher paying firms. Figure 9 complements these results by plotting estimates of $Corr(x_i, \psi_{j(i,t)})$ and $2Cov(x_i, \psi_{j(i,t)})$ for counterfactual economies with different values of s. These findings indicate that production complementarities are the driving force of the strong positive correlation between worker and firm effects and the significant inequality contribution from worker sorting.

6.5 Progressive taxation and allocative efficiency

As shown in Section 4.4, the government can improve the allocation of workers to firms in two ways. First, a less progressive tax system may reduce the misallocation that arise from the tax wedge. Second, letting the tax rates vary across markets may improve allocation by counteracting the differences in the wage setting power of firms. We now use the estimated model to perform a counterfactual that quantifies the impacts of such a tax reform on the



Figure 9: Worker sorting with counterfactual values of $g_j(x)$ and θ_j

Notes: In this figure, we reduce the heterogeneity across firms in amenities or production complementarities by replacing either $g_j(x)$ with $(1-s)g_j(x) + s\bar{g}_j$ or θ_j with $(1-s)\theta_j + s\bar{\theta}$, where $\bar{g}_j = \mathbb{E}_x [g_j(x)], \bar{\theta} = \mathbb{E}[\theta_j]$. Here, $s \in [0, 1]$ is the shrink rate with s = 0 corresponding to the baseline model. We report the share of log earnings variance explained by sorting (subfigure a) and the sorting correlation (subfigure b).

equilibrium allocation and outcomes, including wages, output and welfare.

The counterfactual we consider involves two changes to the monopsonistic labor market. First, we eliminate the tax wedge in the first order condition, which distorts the worker's ranking of firms in favor of those with better amenities. This is done by setting the tax progressivity $(1 - \lambda)$ equal to zero. Second, we remove the labor wedges in the first order conditions of the firms. These wedges causes misallocation of workers across firms with different degree of wage setting power. As proven in Online Appendix C.8, the labor wedges can be eliminated by setting τ_r equal to the labor wedge $1 + \frac{\rho_r}{\lambda\beta}$ in each market r. After changing these parameters of the model, we solve for the new equilibrium allocation and outcomes, including wages, output and welfare. For a set of wages $\{W_{jt}(X, V)\}_{j,t}$ and a tax policy (λ, τ) , we define the welfare as:

$$\mathbb{W}_t = \mathbb{E}\left[\max_{j} u_{it} \left(j, (1+\phi_t)\tau W_{jt}(X_i, V_{it})^{\lambda}\right)\right]$$

where ϕ_t is set so that profits and tax revenues are distributed among all the workers in proportion to their earnings:

$$\underbrace{\phi_t \cdot \mathbb{E}\left[\tau W_{jt}(X_i, V_i)^{\lambda}\right]}_{\text{redistribution}} = \underbrace{\frac{1}{N} \sum_{\text{profits}} \Pi_{jt}}_{\text{profits}} + \underbrace{\mathbb{E}\left[W_j(X_i, V_{it}) - \tau W_j(X_i, V_{it})^{\lambda}\right]}_{\text{government revenue}}$$

In other words, we redistribute profits and tax revenues in a non-distortionary way.

The result are presented in Table 9. They suggest the monopsonistic labor market creates significant misallocation of workers to firms. Eliminating labor and tax wedges increase total welfare by 5 percent and total output by 3 percent. We also find that removing these wedges would increase the sorting of better workers to higher paying firms and lower the rents that workers earn from ongoing employment relationships.

		(1) Monopsonistic Labor Market	(2) No Labor or Tax Wedges	Difference between (1) and (2)
Log of Expected Output	$\log \mathbb{E}[Y_{jt}]$	11.38	11.41	0.03
Total Welfare (log dollars)		12.16	12.21	0.05
Sorting Correlation	$Cor(\psi_{jt}, x_i)$	0.44	0.47	0.03
Labor Wedges	$1 + \frac{\rho_r}{\beta\lambda}$	1.15	1.00	-0.15
Worker Rents (as share of earnings):	P			
Firm-level	$\frac{\rho_r}{\rho_r + \beta \lambda}$	13.3%	12.3%	-1.0%
Market-level	$\frac{\frac{1}{1+\beta\lambda}}{1+\beta\lambda}$	18.0%	16.7%	-1.3%

Table 9: Consequences for Worker Allocation and Outcomes of Eliminating Tax and Labor Wedges

Notes: This table compares the monopsonistic labor market to a counterfactual economy which differs in two ways. First, we eliminate the tax wedge in the first order condition by setting the tax progressivity $(1 - \lambda)$ equal to zero. Second, we remove the labor wedges in the first order conditions of the firms by setting τ_r equal to the labor wedge $1 + \frac{\rho_r}{\lambda\beta}$ in each market r. After changing these parameters of the model, we solve for the new equilibrium allocation and outcomes, including wages, output and welfare. Results are displayed for output, welfare, the sorting correlation, the mean labor wedge, and worker rents.

7 Conclusion

The goal of our paper was to quantify the importance of imperfect competition in the U.S. labor market by estimating the size of rents earned by American firms and workers from ongoing employment relationships. To this end, we constructed a matched employer-employee panel data set by combining the universe of U.S. business and worker tax records for the period 2001-2015. Using this panel data, we described several important features of the U.S. labor market, including the magnitudes of firm-specific wage premiums, the sorting of workers to firms, the production complementarities between high ability workers and productive firms, and the pass through of firm and market shocks to workers' wages. Guided by these empirical results, we developed, identified and estimated a model of the labor market that allowed us to draw inference about imperfect competition, compensating differentials and rent sharing. We also used the model to quantify the relevance of non-wage job characteristics and imperfect competition for inequality and tax policy, to assess the economic determinants of worker sorting, and to offer a unifying explanation of key empirical features of the U.S. labor market.

When considering the interpretation and generality of our study, we emphasize a few caveats. First, we focus on distortions in the allocation of workers to firms and markets. However, tax and labor wedges may also distort the choices of whether and how much to work. Second, our structural model makes several simplifying assumptions, partly to prove identification and take the model to the data. For example, we model individual behavior, and hence do not consider any interdependencies between spouses in the choices of whether and where to work.¹⁴ Moreover, we assume no mobility costs or search frictions, and we do not explicitly model human

 $^{^{14}}$ Autor et al. (2019) and Blundell et al. (2016) estimate a life cycle model with two earners jointly making consumption and labor supply decisions. Their findings suggest an important role for consumption smoothing through household labor supply.

capital investments or wage growth over the life cycle. Third, ideally one would observe and be able to exogenously change both the component of productivity that is specific to a firm and the component of productivity that is common to all firms in such markets. Since we do not have such an ideal experiment, we instead have to reply on the panel data. We show the assumptions under which these data can be used to identify the parameters of interest, and we discuss how some – but not all – these assumptions may be relaxed if one have access to instruments for productivity at both the firm and the market level. Fourth, we consider an equilibrium where each firm views itself as infinitesimal within the market. This assumption is motivated by the fact that very few firms in the U.S. have a large share of the local labor market (as measured by commuting zone). Thus, optimizing firms would essentially ignore the negligible effect of changing their own wages on the overall supply of workers to the market as a whole.¹⁵ For these (and other) reasons, our estimates should not be interpreted as a full accounting of the causes and consequences of imperfect competition in the U.S. labor market.

¹⁵Atkeson and Burstein (2008) relax this assumption in the context of the product market. Berger et al. (2019) adapts such an approach to modeling imperfect competition in the labor market, adding decreasing marginal returns, capital and general equilibrium.

References

- ABOWD, J. A. AND T. LEMIEUX (1993): "The effects of product market competition on collective bargaining agreements: The case of foreign competition in Canada," *The Quarterly Journal of Economics*, 108, 983–1014.
- ABOWD, J. M., F. KRAMARZ, AND D. N. MARGOLIS (1999): "High Wage Workers and High Wage Firms," *Econometrica*, 67, 251–333.
- AKERMAN, A., I. GAARDER, AND M. MOGSTAD (2015): "The skill complementarity of broadband internet," *The Quarterly Journal of Economics*, 130, 1781–1824.
- ANDREWS, M. J., L. GILL, T. SCHANK, AND R. UPWARD (2008): "High Wage Workers and Low Wage Firms: Negative Assortative Matching or Limited Mobility Bias?" Journal of the Royal Statistical Society. Series A (Statistics in Society), 171, 673–697.
- ATKESON, A. AND A. BURSTEIN (2008): "Pricing-to-market, trade costs, and international relative prices," *American Economic Review*, 98, 1998–2031.
- AUTOR, D., A. KOSTØL, M. MOGSTAD, AND B. SETZLER (2019): "Disability Benefits, Consumption Insurance, and Household Labor Supply," *American Economic Review*, 109.
- BARTH, E., A. BRYSON, J. C. DAVIS, AND R. FREEMAN (2016): "It's where you work: Increases in earnings dispersion across establishments and individuals in the US," *Journal of Labor Economics*, 34, S67–S97.
- BERGER, D. W., K. F. HERKENHOFF, AND S. MONGEY (2019): "Labor Market Power," Working Paper 25719, National Bureau of Economic Research.
- BHASKAR, V., A. MANNING, AND T. TO (2002): "Oligopsony and Monopsonistic Competition in Labor Markets," *The Journal of Economic Perspectives*, 16, 155–174.
- BLUNDELL, R., L. PISTAFERRI, AND I. SAPORTA-EKSTEN (2016): "Consumption inequality and family labor supply," *American Economic Review*, 106, 387–435.
- BOAL, W. M. AND M. R. RANSOM (1997): "Monopsony in the labor market," Journal of Economic Literature, 35, 86–112.
- BONHOMME, S. AND G. JOLIVET (2009): "The pervasive absence of compensating differentials," Journal of Applied Econometrics, 24, 763–795.
- BONHOMME, S., T. LAMADON, AND E. MANRESA (2019): "A Distributional Framework for Matched Employee Employee Data," *Econometrica*, 87, 699–739.
- BOROVICKOVA, K. AND R. SHIMER (2017): "High Wage Workers Work for High Wage Firms," Working Paper 24074, National Bureau of Economic Research.

- CAHUC, P. AND G. LAROQUE (2014): "Optimal taxation and monopsonistic labor market: Does monopsony justify the minimum wage?" *Journal of Public Economic Theory*, 16, 259–273.
- CARD, D., A. R. CARDOSO, J. HEINING, AND P. KLINE (2018): "Firms and Labor Market Inequality: Evidence and Some Theory," *Journal of Labor Economics*, 36, S13–S70.
- CARD, D., F. DEVICIENTI, AND A. MAIDA (2013a): "Rent-sharing, holdup, and wages: Evidence from matched panel data," *The Review of Economic Studies*, 81, 84–111.
- CARD, D., J. HEINING, AND P. KLINE (2013b): "Workplace heterogeneity and the rise of West German wage inequality," The Quarterly Journal of Economics, 128, 967–1015.
- CARLSSON, M., J. MESSINA, AND O. N. SKANS (2016): "Wage adjustment and productivity shocks," *The Economic Journal*, 126, 1739–1773.
- DUNNE, T., L. FOSTER, J. HALTIWANGER, AND K. R. TROSKE (2004): "Wage and productivity dispersion in United States manufacturing: The role of computer investment," *Journal* of Labor Economics, 22, 397–429.
- EECKHOUT, J. AND P. KIRCHER (2011): "Identifying Sorting In Theory," *The Review of Economic Studies*, 78, 872–906.
- EKELAND, I., J. J. HECKMAN, AND L. NESHEIM (2004): "Identification and estimation of hedonic models," *Journal of Political Economy*, 112, S60–S109.
- FAGGIO, G., K. G. SALVANES, AND J. VAN REENEN (2010): "The evolution of inequality in productivity and wages: panel data evidence," *Industrial and Corporate Change*, 19, 1919– 1951.
- FRIEDRICH, B., L. LAUN, C. MEGHIR, AND L. PISTAFERRI (2019): "Earnings Dynamics and Firm-Level Shocks," Working Paper 25786, National Bureau of Economic Research.
- GARIN, A. AND F. SILVÉRIO (2019): "How Responsive are Wages to Demand within the Firm? Evidence from Idiosyncratic Export Demand Shocks," *Working paper*.
- GIBBONS, R. AND L. KATZ (1992): "Does unmeasured ability explain inter-industry wage differentials?" *The Review of Economic Studies*, 59, 515–535.
- GIBBONS, R., L. F. KATZ, T. LEMIEUX, AND D. PARENT (2005): "Comparative advantage, learning, and sectoral wage determination," *Journal of Labor Economics*, 23, 681–724.
- GUISO, L., L. PISTAFERRI, AND F. SCHIVARDI (2005): "Insurance within the firm," Journal of Political Economy, 113, 1054–1087.
- GUNER, N., R. KAYGUSUZ, AND G. VENTURA (2014): "Income taxation of US households: Facts and parametric estimates," *Review of Economic Dynamics*, 17, 559–581.
- HAMERMESH, D. S. (1999): "Changing inequality in markets for workplace amenities," *The Quarterly Journal of Economics*, 114, 1085–1123.

- HEATHCOTE, J., K. STORESLETTEN, AND G. L. VIOLANTE (2017): "Optimal tax progressivity: An analytical framework," *The Quarterly Journal of Economics*, 132, 1693–1754.
- KLINE, P., N. PETKOVA, H. WILLIAMS, AND O. ZIDAR (2018a): "Who Profits from Patents? Rent-Sharing at Innovative Firms," Working Paper 25245, National Bureau of Economic Research.
- KLINE, P., R. SAGGIO, AND M. SØLVSTEN (2018b): "Leave-out estimation of variance components," *Working paper*.
- LAMADON, T. (2016): "Productivity Shocks, Long-Term Contracts and Earnings Dynamics," Working paper.
- LASSERRE, J. B. (2001): "Global optimization with polynomials and the problem of moments," SIAM Journal on Optimization, 11, 796–817.
- LAVETTI, K. AND I. M. SCHMUTTE (2017): "Estimating compensating wage differentials with endogenous job mobility," *Working paper*.
- MAESTAS, N., K. J. MULLEN, D. POWELL, T. VON WACHTER, AND J. B. WENGER (2018):"The Value of Working Conditions in the United States and Implications for the Structure of Wages," Working Paper 25204, National Bureau of Economic Research.
- MANNING, A. (2003): Monopsony in motion: Imperfect competition in labor markets, Princeton University Press.
- (2011): "Imperfect competition in the labor market," *Handbook of Labor Economics*, 4, 973–1041.
- MARINESCU, I. AND R. RATHELOT (2018): "Mismatch Unemployment and the Geography of Job Search," *American Economic Journal: Macroeconomics*, 10, 42–70.
- MAS, A. AND A. PALLAIS (2017): "Valuing alternative work arrangements," American Economic Review, 107, 3722–59.
- MIRRLEES, J. A. (1971): "An exploration in the theory of optimum income taxation," *The Review of Economic Studies*, 38, 175–208.
- MORTENSEN, D. (2003): Wage dispersion: why are similar workers paid differently?, MIT Press.
- MURPHY, K. M. AND R. H. TOPEL (1990): "Efficiency wages reconsidered: Theory and evidence," in Advances in the Theory and Measurement of Unemployment, Springer, 204–240.
- PIERCE, B. (2001): "Compensation Inequality," The Quarterly Journal of Economics, 116, 1493–1525.
- ROSEN, S. (1986): "The Theory of Equalizing Differences," *Handbook of Labor Economics*, 1, 641–692.

- SHIMER, R. AND L. SMITH (2000): "Assortative Matching and Search," *Econometrica*, 68, 343–369.
- SONG, J., D. J. PRICE, F. GUVENEN, N. BLOOM, AND T. VON WACHTER (2018): "Firming Up Inequality," *The Quarterly Journal of Economics*, 134, 1–50.
- SORKIN, I. (2018): "Ranking Firms using Revealed Preference," The Quarterly Journal of Economics, 133, 1331–1393.
- TABER, C. AND R. VEJLIN (2016): "Estimation of a Roy/search/compensating differential model of the labor market," *Working paper*.
- VAN REENEN, J. (1996): "The creation and capture of rents: wages and innovation in a panel of UK companies," *The Quarterly Journal of Economics*, 111, 195–226.
- WISWALL, M. AND B. ZAFAR (2017): "Preference for the workplace, investment in human capital, and gender," *The Quarterly Journal of Economics*, 133, 457–507.
- YAGAN, D. (2015): "Capital tax reform and the real economy: The effects of the 2003 dividend tax cut," *American Economic Review*, 105, 3531–3563.
- YAZICI, H. AND C. SLEET (2017): "Taxation, Redistribution and Frictional Labor Supply," in 2017 Meeting Papers, Society for Economic Dynamics, 298.