

Firms and Labor Market Inequality: Evidence and Some Theory

David Card, *Institute for Economic Analysis (Consejo Superior de Investigaciones Científicas) and Barcelona Graduate School of Economics*

Ana Rute Cardoso, *Institute for Employment Research (IAB)*

Joerg Heining, *University of California, Berkeley*

Patrick Kline, *University of California, Berkeley*

We synthesize two related literatures on firm-level drivers of wage inequality. Studies of rent sharing that use matched worker-firm data find elasticities of wages with respect to value added per worker in the range of 0.05–0.15. Studies of wage determination with worker and firm fixed effects typically find that firm-specific premiums explain 20% of overall wage variation. To interpret these findings, we develop a model of wage setting in which workers have idiosyncratic tastes for different workplaces. Simple versions of this model can rationalize standard fixed effects specifications and also match the typical rent-sharing elasticities in the literature.

I. Introduction

How much does where you work determine what you earn? In the standard competitive labor market model, firms take market wages as given, and firm-specific heterogeneity influences who is hired but not the level of pay

We are extremely grateful to Raffaele Saggio for assistance in preparing this paper and to Katherine Shaw and David Green for helpful suggestions on an early draft. Ana Rute Cardoso acknowledges financial support from the Spanish Ministry of

[*Journal of Labor Economics*, 2018, vol. 36, no. S1]

© 2018 by The University of Chicago. All rights reserved. 0734-306X/2018/36S1-0009\$10.00

Submitted March 14, 2016; Accepted June 26, 2017

of any particular worker. The pervasive influence of this perspective is evident in major reviews of the wage inequality literature (Katz and Autor 1999; Goldin and Katz 2009; Acemoglu and Autor 2011), which focuses almost exclusively on the role of market-level skill prices in driving inequality trends.¹ This view stands in stark contrast to the industrial organization literature, which typically models markets as imperfectly competitive (Tirole 1988; Pakes 2016). Although economists seem to agree that part of the variation in the prices of cars and breakfast cereal is due to factors other than marginal cost, the notion that wages differ substantially among equally skilled workers remains highly controversial.

A long tradition in labor economics posits that employers have significant latitude to set wages (e.g., Robinson 1933; Lester 1946; Reynolds 1946; Slichter 1950), a view that found early support in empirical studies of industry wage differentials (Katz 1986; Krueger and Summers 1988; Katz and Summers 1989). Yet it has proven difficult to convincingly distinguish industry wage premia from subtle forms of dynamic sorting (Murphy and Topel 1990; Gibbons and Katz 1992; Gibbons et al. 2005). The growing availability of matched employer-employee data sets offers opportunities to study these questions at the more granular level of the firm. Nevertheless, many of the fundamental identification problems that plagued earlier studies carry over to these new data sets. This review summarizes what has been learned so far from these new data sets about the importance of firms in wage setting and what challenges remain.

Our starting point is the widely accepted finding that observably similar firms exhibit massive heterogeneity in measured productivity (e.g., Syverson 2011). A natural question is whether some of these productivity differences spill over to wages. The *prima facie* case for such a link seems quite strong: a number of recent studies find that trends in aggregate wage dispersion closely track trends in the dispersion of productivity across workplaces (Dunne et al. 2004; Faggio, Salvanes, and Van Reenen 2010; Barth et al. 2016). However, these aggregate relationships are potentially driven in part by changes in the degree to which different groups of workers are assigned to different firms.

Two distinct literatures attempt to circumvent the sorting issue using linked employer-employee data. The first literature studies the impact of shocks to firm productivity on the wages of workers. The resulting esti-

Economy and Competitiveness (Severo Ochoa Programme for Centres of Excellence in Research and Development grant SEV-2015-0563) and the Research Council of Norway (Europe in Transition funding scheme project 227072/F10 at the Centre for the Study of Equality, Social Organization, and Performance). Contact the corresponding author, David Card, at card@econ.berkeley.edu. Information concerning access to the data used in this paper is available as supplementary material online.

¹ This market-wide perspective is also common in economic models of discrimination, which typically have no role for firm-specific factors to affect the wages of female or minority workers (see, e.g., Charles and Guryan 2008, 2011).

mates are typically expressed as rent-sharing elasticities. A review of this literature suggests that estimated rent-sharing elasticities are surprisingly robust to the choice of productivity measure and labor market environment: most studies that control for worker heterogeneity find wage-productivity elasticities in the range of 0.05–0.15, although a few older studies find larger elasticities. We also provide some new evidence on the relationship between wages and firm-specific productivity using matched worker-firm data from Portugal. We investigate a number of specification issues that frequently arise in this literature, including the impact of filtering out industry-wide shocks, different approaches to measuring rents, and econometric techniques for dealing with unobserved worker heterogeneity.

A second literature that builds on the additive worker and firm effects wage model proposed by Abowd, Kramarz, and Margolis (1999; hereafter, AKM) uses data on wage outcomes as workers move between firms to estimate firm-specific pay premiums. This literature also finds that firms play an important role in wage determination, with a typical finding that about 20% of the variance of wages is attributable to stable firm wage effects. We discuss some of the issues that arise in implementing AKM's two-way fixed effects estimator, which is the main tool used in this literature, and evidence on the validity of the assumptions underlying the AKM specification.

We then attempt to forge a more direct link between the rent-sharing literature and studies based on the AKM framework. Using data from Portugal we show that more productive firms pay higher average wage premiums but also tend to hire more productive workers. Indeed, we estimate that about 40% of the observed difference in average hourly wages between more and less productive firms is attributable to the differential sorting of higher-ability workers to more productive firms, underscoring the importance of controlling for worker heterogeneity. Consistent with the additive specification underlying the AKM model, we find that the wage premiums offered by more productive firms to more or less educated workers are very similar and that the relative wage of highly educated workers is remarkably stable across firms.

In the final section of the paper we develop a stylized model of imperfect competition in the labor market that provides a tractable framework for studying the implications of worker and firm heterogeneity for wage inequality. Our analysis builds on the static partial equilibrium monopsony framework introduced by Joan Robinson (1933), which, as noted by Manning (2011), captures many of the same economic forces as search models, albeit without providing a theory of worker flows between jobs. We provide a microeconomic foundation for imperfect labor market competition by allowing workers to have heterogeneous preferences over the work environments of different potential employers.² This workplace differentiation can reflect heterogeneity in firm location, job characteristics (e.g., cor-

² In their review of monopsony models, Boal and Ransom (1997) refer to this as the case of classic differentiation.

porate culture, starting times for work), or other factors that are valued differently by different workers. Regardless of its source, such heterogeneity makes employers imperfect substitutes in the eyes of workers, which in turn gives firms some wage-setting power.

To capture this heterogeneity, we adopt a random utility model of worker preferences that characterizes firm-specific labor supply functions. We presume, as in Robinson's analysis and much of the industrial organization literature, that the firm cannot price discriminate on the basis of a worker's idiosyncratic preference for the firm's work environment. Hence, rather than offer each worker her reservation wage (e.g., as in Postel-Vinay and Robin 2002), firms post a common wage for each skill group that is marked down from marginal product in inverse proportion to their elasticity of labor supply to the firm. This condition provides a natural analog to the equilibrium markups of price over marginal cost found in workhorse models of differentiated product demand and pricing (Berry 1994; Berry, Levinsohn, and Pakes 1995; Pakes 2016).

We show that many well-documented empirical regularities can be rationalized in this framework. Firm heterogeneity in productivity affects not only the firm size distribution but also the distribution of firm-specific wage premiums and the degree of sorting of different skill groups across firms. Calibrating a simple version of the model to observed rent-sharing estimates yields predicted wage dispersion in excess of what has been found in a wide class of search models (Hornstein, Krusell, and Violante 2011). We also present conditions under which log wages are additively separable into worker and firm components, as in the pioneering econometric model of AKM. Specifically, we show that the firm-specific wage premium will be constant across skill groups if they are perfect substitutes in production and if different skill groups have a similar relative valuation for wages versus nonwage components in their assessment of alternative jobs. Even under these conditions, however, the market-level wage gap between skill groups will reflect differences in their employment distributions across more and less productive firms. More generally, groups that put a higher relative value on wages in comparing different jobs will receive a larger wage premium at more profitable firms.

We conclude with some thoughts on unresolved empirical and theoretical issues in the literature. An important challenge for empirical work on rent sharing is the identification of firm-specific shocks to productivity, because correlated shocks can yield market-level wage adjustments. In the firm-switching literature, a key question is whether conventionally estimated firm-specific pay premiums predict the wage changes associated with exogenously induced job accessions and separations. On the theoretical side, important questions include how to model strategic labor market interactions between firms and to what extent the insights from simple static wage-setting models of workplace differentiation carry over to dynamic labor market settings with search or mobility frictions.

II. Productivity, Wages, and Rent Sharing

A large empirical literature reviewed by Syverson (2011) documents that firms, like workers, exhibit vast heterogeneity in productivity. For example, Syverson (2004) finds that the 90th and 10th percentiles of total factor productivity (TFP) among US manufacturing firms differ by an average factor of approximately 2 within four-digit industries. Hsieh and Klenow (2009) find even larger productivity gaps in India and China, with 90/10 TFP ratios on the order of 5. While the variation in measured productivity probably overstates the true heterogeneity in plant-level efficiency, there is also strong evidence in the literature that measured productivity conveys real information. For example, measured TFP is a strong predictor of firm survival (Foster, Haltiwanger, and Syverson 2008).

It is natural to wonder whether these large productivity differences lead to differences in worker pay. In fact, an extensive literature has documented the existence of substantial wage differences across plants and establishments (Lester 1946; Slichter 1950; Davis and Haltiwanger 1991; Groshen 1991; Bernard and Jensen 1995; Cardoso 1997, 1999; Skans, Edin, and Holmlund 2009; Song et al. 2015) that are strongly correlated with basic measures of productivity. Nevertheless, economists have been reluctant to interpret these differences as wage premiums or rents, since it has been difficult to know how unobserved worker quality differs across plants.

Recent studies, however, have documented some striking links between establishment-level productivity and wage dispersion (Dunne et al. 2004; Faggio et al. 2010; Barth et al. 2016). Figure 1 plots results from Barth et al. (2016), showing remarkably similar trends in the dispersion of wages and productivity across business establishments in the United States. Taken at

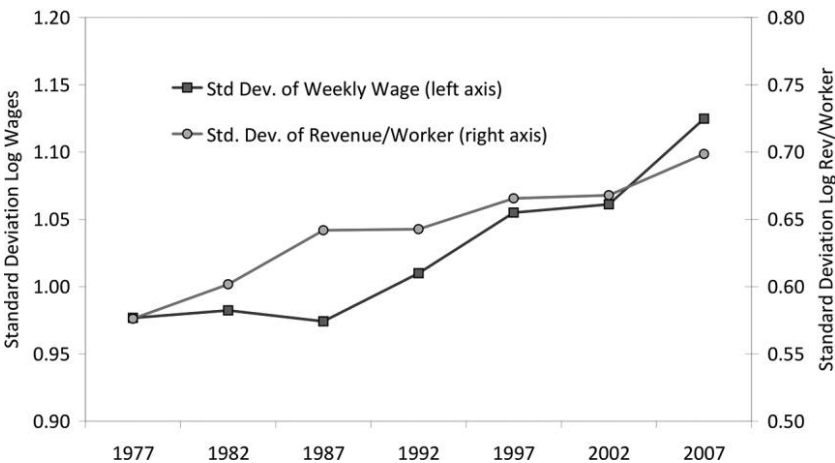


FIG. 1.—Trends in between-establishment dispersion in wages and productivity. Source: Barth et al. (2016). A color version of this figure is available online.

face value, these parallel trends are consistent with a roughly unit elasticity of establishment wages with respect to productivity (see Barth et al. 2016, S71). Of course, figure 1 does not tell us whether the composition of the workforce employed at these establishments is changing over time. What appear to be more productive establishments may simply be establishments that hire more skilled workers, which is fully consistent with the competitive labor market model in which all firms pay the same wages for any given worker.

A more direct attack on the question of whether firm-specific productivity differentials feed into differences in wages comes from the empirical literature on rent sharing. Table A1 describes 22 recent studies in this literature. The basic idea in these papers is to relate wages to some measure of employer profitability or rents. This is, in many respects, the analog of a large literature in international economics and industrial organization on the pass-through of cost shocks to prices. A typical finding in that literature is that shocks to marginal cost (typically measured via exchange rate fluctuations) yield muted fluctuations in product prices, signaling the presence of market power (Goldberg and Hellerstein 2013; Gorodnichenko and Talavera 2017). Most rent-sharing studies find similarly imperfect propagation of productivity shocks to wages. However, the elasticities reported in this literature are often sensitive to measurement issues, which we now review.

A. Measuring Rents

The empirical rent-sharing literature is motivated by an assumed structural relationship between wages and either profit per worker or a measure of quasi rent per worker. To facilitate discussion, suppose that there is a single type of labor at a firm j and that the wage (w_j) is determined by a structural relationship of the form

$$w_j = b + \frac{\gamma Q_j}{N_j}, \quad (1)$$

where b represents an alternative wage, N_j is employment at the firm, Q_j represents quasi rents, and γ is a rent-sharing parameter.³ The firm combines labor inputs and capital (K_j) and faces an exogenous rental rate r on capital, yielding the quasi rent

³ This equation can be derived by assuming that workers as a group bargain with the firm to maximize a weighted product of profits and the excess wage bill $N_j(w_j - b)$, where the weight on the wage bill is γ . See, e.g., de Menil (1971). Strictly speaking, this derivation assumes that Q/N is exogenous to the level of wages, as is the case when employment and capital are efficiently determined by joint bargaining. See Svejnar (1986) and Card, Devicienti, and Maida (2014) for a related discussion of potential holdup problems in the determination of capital.

$$Q_j = VA_j - bN_j - rK_j,$$

where VA_j is value added (revenue net of materials costs). Value added is related to labor and capital inputs by

$$VA_j = P_j T_j f(N_j, K_j),$$

where P_j is a potentially firm-specific selling price index, T_j is an index of technical efficiency, and f is a standard production function. Here $P_j T_j$ represents total factor productivity (TFP_j), which, in the terminology of Foster et al. (2008), is also referred to as revenue productivity because it is the product of physical productivity T_j and product price P_j .

We assume that TFP_j is the driving source of variation that researchers are implicitly trying to model in the rent-sharing literature. Under this interpretation, firm-specific TFP shocks lead to changes in quasi rent per worker that cause wages to fall or rise relative to the alternative wage. The elasticity of wages with respect to an exogenous change in quasi rent per worker is

$$\xi_{Q_j} = \frac{\gamma(Q_j/N_j)}{b_j + \gamma(Q_j/N_j)}, \quad (2)$$

which corresponds to the share of rents in wages. The elasticity of wages with respect to profit per worker (π_j/N_j) should be of comparable magnitude. Indeed, under the usual bargaining interpretation of equation (1), profits per worker are a constant share of quasi rents per worker:

$$\frac{\pi_j}{N_j} = (1 - \gamma) \frac{Q_j}{N_j}.$$

Rather than measure quasi rents, a majority of studies relate wages to value added per worker. The elasticity of wages with respect to value added per worker is

$$\xi_j = \xi_{Q_j} \times \frac{VA_j}{Q_j},$$

which will be bigger than ξ_{Q_j} , since $Q_j < VA_j$.⁴ For example, data reported by Card et al. (2014) suggest that the ratio of value added to quasi rent for firms in Northeast Italy is typically around 2. Finally, some studies in the literature relate wages to revenue per worker rather than to value added per worker. Under the assumption that intermediate input costs represent a constant share of revenues, the elasticity of wages with respect to revenue

⁴ Again, this is true only under the assumption that the levels of value added and quasi rents are exogenously determined, as in an efficient bargaining model. More generally, the ratio of value added per worker to quasi rent per worker can depend on the wage.

per worker will be equal to the elasticity with respect to value added per worker. More generally, if intermediate input costs vary—for example, because of varying energy costs—and these costs are passed through to the firm's consumers, one would expect a smaller elasticity of wages with respect to revenues per worker.

An important confounding factor in the rent-sharing literature is variation in worker quality. Firms that employ more highly skilled workers would be expected to have higher revenue per worker and value added per worker and also pay higher wages, leading to a potential upward bias in the measured rent-sharing elasticity in cross-sectional studies that compare different firms at a point in time. A similar bias can also arise in longitudinal studies that compare changes in firm-specific wages and profitability over time if there are unobserved changes in the skill characteristics of workers. For this reason, when we summarize the studies in the literature, we classify studies based on whether the research design includes controls for unobserved worker skills.

B. A Summary of the Rent-Sharing Literature

Table 1 synthesizes the estimated rent-sharing elasticities from the 22 studies listed in table A1, extracting one or two preferred specifications from each study and adjusting all elasticities to an approximate value added per worker basis.⁵ We divide the studies into three broad generations on the basis of the level of aggregation in the measures of rents and wages.

The first group of studies, which includes two influential papers from the early 1990s, uses industry-wide measures of productivity and either individual-level or firm-wide average wages. The average rent-sharing elasticity in this group is 0.16. A second generation of studies includes five papers, mostly from the mid-1990s, that use firm- or establishment-specific measures of rents but measure average wages of employees at the workplace level. The average rent-sharing elasticity in this group is 0.15, although there is a relatively wide range of variation across the studies. Given the likely problems caused by variation in worker quality, we suspect that most first-generation and second-generation studies yield upward-biased estimates of the rent-sharing elasticity.

A third generation of studies consists of 18 relatively recent papers that study the link between firm- or establishment-specific measures of rents and individual-specific wages. Many of these studies attempt to control for variation in worker quality, in some cases by studying the effect of changes in measured rents on changes in wages. In this group the mean rent-sharing elasticity is 0.08, although a few studies report rent-sharing

⁵ We extract an IV estimate when one is available and convert elasticities with respect to profit per worker or quasi rent per worker to a value added per worker basis by multiplying by 2.

Table 1
Summary of Estimated Rent-Sharing Elasticities from the Recent Literature
(Preferred Specification, Adjusted to Total Factor Productivity Basis)

Study	Country/Industry	Estimated Elasticity	Standard Error
Group 1—Industry-level profit measure:			
Christofides and Oswald 1992	Canadian manufacturing	.140	.035
Blanchflower, Oswald, and Sanfey 1996	US manufacturing	.060	.024
Esteveao and Tevlin 2003	US manufacturing	.290	.100
Group 2—Firm-level profit measure, mean firm wage:			
Abowd and Lemieux 1993	Canadian manufacturing	.220	.081
Van Reenen 1996	UK manufacturing	.290	.089
Hildreth and Oswald 1997	United Kingdom	.040	.010
Hildreth 1998	UK manufacturing	.030	.010
Barth et al. 2016	United States	.160	.002
Group 3—Firm-level profit measure, individual-specific wage:			
Margolis and Salvanes 2001	French manufacturing	.062	.041
Margolis and Salvanes 2001	Norwegian manufacturing	.024	.006
Arai 2003	Sweden	.020	.004
Guiso et al. 2005	Italy	.069	.025
Fakhfakh and FitzRoy 2004	French manufacturing	.120	.045
Du Caju et al. 2011	Belgium	.080	.010
Martins 2009	Portuguese manufacturing	.039	.021
Gürtzgen 2009	Germany	.048	.002
Cardoso and Portela 2009	Portugal	.092	.045
Arai and Heyman 2009	Sweden	.068	.002
Card et al. 2014	Italy (Veneto region)	.073	.031
Carlsson et al. 2014	Swedish manufacturing	.149	.057
Card et al. 2016	Portugal, between firm	.156	.006
Card et al. 2016	Portugal, within job	.049	.007
Bagger et al. 2014	Danish manufacturing	.090	.020

NOTE.—For a more complete description of each study, see table A1.

elasticities that are 0.05 or smaller. To quantify the potential role of rent sharing in the overall dispersion of wages, note that the standard deviation of average value added per worker across firms in the Portuguese data that we analyze in the next section is about 0.70, with a spread between the 90th and 10th percentiles of around 1.6. An elasticity of 0.08 implies that the variation in productivity across firms creates a Lester range of wage variability (Lester 1952) of about 13 log points between firms at the 90th and 10th percentiles.

Although significant progress has been made in this literature, none of these studies is entirely satisfactory. Very few studies have clear exogenous sources of variation in productivity. Most papers (e.g., Guiso, Pistaferri, and Schivardi 2005; Carlsson, Messina, and Skans 2014; Card, Cardoso, and

Kline 2016) rely on timing assumptions about the stochastic process driving productivity to justify using lags as instruments. A notable exception is Van Reenen (1996), who studies the effects of major firm innovations on employee wages. He finds a very large rent-sharing elasticity of 0.29, but this figure may be upward biased by skill upgrading on the part of innovative firms, a concern he could not address with aggregate data. Other studies (e.g., Abowd and Lemieux 1993; Card et al. 2014) use industry-level shocks as instruments for productivity. However, these instruments may violate the exclusion restriction if labor supply to the sector is inelastic, since even fully competitive models predict that industry-level shocks can yield equilibrium wage responses. Moreover, industry-level shocks might yield general equilibrium responses that change worker's outside options (Beaudry, Green, and Sand 2012). Finally, with the move to matched employer-employee micro data, economists have had to contend with serious measurement error problems that emerge when constructing plant-level productivity measures. It remains to be seen whether instrumenting using lags fully resolves these issues.

C. Specification Issues: A Replication in Portuguese Data

To supplement the estimates in the literature and probe the impact of different design choices on the magnitude of the resulting elasticities, we conducted our own analysis of rent-sharing effects using matched employer-employee data from Portugal. The wage data for this exercise come from Quadros de Pessoal (QP), a census of private sector employees conducted each October by the Portuguese Ministry of Employment. We merge these data to firm-specific financial information from the Sistema de Analisis de Balances Ibericos (SABI) database, distributed by Bureau van Dijk.⁶ We select all male employees observed between 2005 and 2009 who work in a given year at a firm in the SABI database with valid information on sales per worker for each year from 2004 to 2010 and on value added per worker for each year from 2005 to 2009.

Panel A of table 2 presents a series of specifications in which we relate the log hourly wage observed for a worker in a given year (between 2005 and 2009) to mean log value added per worker or mean log sales per worker at his employer, averaged over the sample period. These are simple cross-sectional rent-sharing models in which we use an averaged measure of rents at the employer to smooth out the transitory fluctuations and measurement

⁶ Businesses in Portugal are required to file income statements and balance sheet information annually at their local Commercial Registry (Conservatoria do Registo Comercial). These reports are publicly accessible and are collected by financial service firms and assembled into the SABI database. We merge SABI and QP using information on detailed location, industry, firm creation date, shareholder equity, and annual sales that are available in both data sets. See Card et al. (2016) for more information on the matching process.

Table 2
Cross-Sectional and Within-Job Models of Rent Sharing for Portuguese Male Workers

	Basic Specification (1)	Basic + Major Industry/City (2)	Basic + Detailed Industry/City (3)
A. Cross-sectional models (worker-year observations, 2005–9):			
OLS: rent measure = mean log value added per worker, 2005–9	.270 (.017)	.241 (.015)	.207 (.011)
OLS: rent measure = mean log sales per worker, 2005–9	.153 (.009)	.171 (.007)	.159 (.004)
IV: rent measure = mean log value added per worker, 2005–9; instrument = mean log sales per worker, 2004–10	.327 (.014)	.324 (.011)	.292 (.008)
First-stage coefficient	.475 (<i>t</i> = 26.19)	.541 (<i>t</i> = 40.72)	.562 (<i>t</i> = 64.38)
B. Within-job models (change in wages from 2005 to 2009 for stayers):			
OLS: rent measure = change in log value added per worker from 2005 to 2009	.041 (.006)	.039 (.005)	.034 (.003)
OLS: rent measure = change in log sales per worker from 2005 to 2009	.015 (.005)	.014 (.004)	.013 (.003)
IV: rent measure = change in log value added per worker from 2005 to 2009; instrument = change in log sales per worker, 2004–10	.061 (.018)	.059 (.017)	.056 (.016)
First-stage coefficient	.221 (<i>t</i> = 11.82)	.217 (<i>t</i> = 13.98)	.209 (<i>t</i> = 18.63)

NOTE.—The sample in panel A is 2,503,336 person-year observations from Quadros de Pessoal (QP) for males working in 2005–9 between the ages of 19 and 65 years with at least 2 years of potential experience employed at a firm with complete value-added data (from Sistema de Analisis de Balances Ibericos [SABI]) for 2005–9 and sales data (from QP) for 2004 and 2010. The sample in panel B is 284,071 males ages 19–61 years in 2005 who worked every year from 2005 to 2009 at a firm with complete value-added data (from SABI) for 2005–9 and sales data (from QP) for 2004 and 2010. Standard errors are clustered by firm (62,845 firms in panel A, 44,661 firms in panel B). Models in panel A control for cubic in experience and unrestricted education × year dummies. Models in panel B control for a quadratic in experience and education. Models in col. 2 also control for 20 major industries and two major cities (Lisbon and Porto). Models in col. 3 also control for 202 detailed industry dummies and 29 Nomenclature of Territorial Units for Statistics region 3 location dummies. IV = instrumental variables; OLS = ordinary least squares.

errors in the financial data. In row 1 we present models using mean log value added per worker as the measure of rents; in row 2 we use mean log sales per worker; and in row 3 we use mean log value added per worker over the 2005–2009 period but instrument this with mean log sales per worker over

a slightly wider window (2004–2010). For each choice we show a basic specification (with only basic human capital controls) in column 1, a richer specification with controls for major industry and city in column 2, and a full specification with dummies for 202 detailed industries and 29 regions in column 3.

Two main conclusions emerge from these simple models. First, the rent-sharing elasticity is systematically larger when rents are measured by value added per worker than by sales per worker.⁷ Second, the rent-sharing elasticities from this approach are relatively large. Interestingly, the 0.20–0.30 range of estimates is comparable to the range of the studies in groups 1 and 2 of table 1.

An obvious concern with the specifications used in panel A is that they fail to fully control for variation in worker quality. As discussed above, this is likely to lead to an upward bias in the relationship between wages and value added per worker. The specifications in panel B of table 2 partially address this by examining the effect of changes in firm-specific rents on changes in wages for workers who remain at the firm over the period from 2005 to 2009, a within-job or stayers design. We present three sets of specifications of this design. The models in row 4 measure the change in rents by the change in log value added per worker. The models in row 5 use the change in log sales per worker. The models in row 6 use the change in value added per worker as the measure of rents but instrument the change using the change in sales per worker over a slightly wider interval to reduce the impact of measurement errors in value added.⁸

Relative to the cross-sectional models, the within-job models yield substantially smaller rent-sharing elasticities. This difference is likely due to some combination of unobserved worker quality in the cross-sectional designs (which leads to an upward bias in these specifications), measurement error (which causes a larger downward bias in the stayer designs), and the fact that value-added fluctuations may include a transitory component that firms insure workers against (Guiso et al. 2005).⁹ The discrepancy is particularly large for ordinary least squares (OLS) models using sales per worker (compare rows 2 and 5 of table 2): the elasticity for stayers is only about one-tenth as large as the cross-sectional elasticity. We suspect that measurement errors and transitory fluctuations in annual sales are relatively large,

⁷ A similar finding is reported by Card et al. (2014) using Italian data.

⁸ If measurement errors in value added per worker in year t are uncorrelated with errors or fluctuations in sales per worker in years $t + 1$ and $t - 1$, then the use of a bracketing instrument will eliminate the effect of measurement error in value added. We suspect that this is only partially true, so the IV approach reduces but does not fully eliminate the effect of errors in value added.

⁹ A third potential explanation is selection bias in the stayer models, induced by selecting a sample of job stayers. Results presented by Card et al. (2016, their table B10) suggest that this factor is relatively small.

and the impact of these factors is substantially magnified in the within-job specifications estimated by OLS. Given the presence of errors and idiosyncratic fluctuations, we prefer the instrumental variables (IV) estimates in row 6, which point toward a rent-sharing elasticity of approximately 0.06.

An interesting feature of both the OLS and the IV within-job estimates is that the addition of detailed industry controls reduces the rent-sharing elasticity by 10%–20%. Since these industry dummies absorb industry-wide productivity shocks that are shared by the firms in the same sector, we conclude that the rent-sharing elasticity with respect to firm-specific productivity shocks (which is estimated by the models in col. 3) is somewhat smaller than the elasticity with respect to sector-wide shocks (which are incorporated in the elasticities in the models in col. 1). If true more generally, this suggests that the use of industry-wide rent measures will lead to somewhat larger rent-sharing elasticities than would be obtained using firm-specific productivity measures and controlling for industry-wide trends. A similar conclusion is reported by Carlsson et al. (2014).

Overall, we conclude that the estimated elasticities of wages with respect to value added per worker in Portugal are comparable in magnitude to the estimates in the existing rent-sharing literature summarized in table 1. A typical estimate in our data and the literature from specifications that control for worker quality is between 0.05 and 0.10, although there are a few estimates on each side of this range.

III. Firm Switching

While the rent-sharing literature documents a strong correlation between firm profitability and pay, a parallel literature finds that workers who move between firms (or establishments) experience wage gains or losses that are highly predictable. In this section we provide an overview of recent findings from this approach and discuss some of the major issues in this literature. In the following section we discuss how the firm-specific wage premiums estimated by studies of firm switching are related to measures of firm profitability, providing a link between the rent-sharing and firm-switching literatures.

A. AKM Models

In their seminal study of the French labor market, Abowd et al. (1999) specified a model for log wages that includes additive effects for workers and firms. Specifically, their model for the log wage of person i in year t takes the form

$$\ln w_{it} = \alpha_i + \psi_{J(i,t)} + X'_{it}\beta + \varepsilon_{it},$$

where X_{it} is a vector of time-varying controls (e.g., year effects and controls for experience), α_i is a person effect capturing the (time-invariant) portable

component of earnings ability, the $\{\psi_j\}_{j=1}^J$ are firm-specific relative pay premiums, $J(i, t)$ is a function indicating the employer of worker i in year t , and ε_{it} is an unobserved time-varying error capturing shocks to human capital, person-specific job match effects, and other factors. The innovation in the AKM framework is the presence of the firm effects, which allow for the possibility that some firms pay systematically higher or lower wages than other firms. Specifically, the AKM model predicts that workers who move from firm k to firm j will experience an average wage change of $\psi_j - \psi_k$, while those who move in the opposite direction will experience an average change of $\psi_k - \psi_j$, a stark symmetry prediction that we discuss in more detail below.

Estimates of AKM-style models on population-level administrative data sets from a variety of different countries have found that the firm effects in these models typically explain 15%–25% of the variance of wages, less than the person effects but enough to indicate that firm-specific wage setting is important for wage inequality.¹⁰ One problem with this assessment is that the person and firm effects are estimated with considerable imprecision, which means that the explanatory power of firms will typically be somewhat overstated, a problem that was also recognized in the earlier literature on industry wage differentials (Krueger and Summers 1988). Andrews et al. (2008) provide an approach to dealing with this problem that we discuss in more detail below.

If different firms pay different wage premiums, the pattern of sorting of workers to firms will also matter for overall wage inequality. In particular, the variance of log wages is

$$\begin{aligned} \text{var}(\ln w_{it}) = & \text{var}(\alpha_i) + \text{var}(\psi_{J(i,t)}) + \text{var}(X'_{it}\beta) + \text{var}(\varepsilon_{it}) \\ & + 2\text{cov}(\alpha_i, \psi_{J(i,t)}) + 2\text{cov}(\alpha_i, X'_{it}\beta) + 2\text{cov}(\psi_{J(i,t)}, X'_{it}\beta), \quad (3) \end{aligned}$$

which includes both the variance of the firm-specific wage premiums and a term reflecting the covariance of the worker and firm effects. If workers with a higher earning capacity are more likely to work at higher-premium firms, then this covariance term will be positive, and any inequality effects from the presence of the firm premiums will be amplified.

¹⁰ For example, Abowd, Lengermann, and McKinney (2003) find that firm effects comprise 17% of the variance of US wages. Card, Heining, and Kline (2013) find that establishment effects explain between 18% and 21% of the variance of the wages of German men, depending on the time period studied. Card et al. (2016) find that firm effects explain 20% of the variance of hourly wages for Portuguese men and 17% of the variance for women. Macis and Schivardi (2016) find that firm effects explain 15% of the wage variance of Italian manufacturing workers. Finally, Lavetti and Schmutte (2016) find that establishment effects explain 21% of the variance of wages of workers in the formal sector in Brazil.

An alternative decomposition uses the fact that

$$\begin{aligned} \text{var}(\ln w_{it}) &= \text{cov}(\ln w_{it}, \alpha_i) + \text{cov}(\ln w_{it}, \psi_{J(i,t)}) \\ &\quad + \text{cov}(\ln w_{it}, X'_{it}\beta) + \text{cov}(\ln w_{it}, \varepsilon_{it}). \end{aligned} \quad (4)$$

This yields an ensemble assessment of the importance of each variance component to wage dispersion that includes the contribution of the covariance between wage components. For example, under this decomposition, the contribution of the firm component to total wage variation would be $\text{cov}(\ln w_{it}, \psi_{J(i,t)}) = \text{var}(\psi_{J(i,t)}) + \text{cov}(\alpha_i, \psi_{J(i,t)}) + \text{cov}(X'_{it}\beta, \psi_{J(i,t)})$. One way to think about this decomposition is that one-half of the firm covariance terms in equation (3) are attributed to the firm-specific wage premiums.

B. Identifying Age and Time Effects

A technical issue that arises with the AKM model is appropriate specification of the effects of age. Following Mincer (1974), it is conventional to include a polynomial in age or potential experience (age minus education minus 6) in X_{it} . However, it is also standard to include a set of year indicators in X_{it} to adjust for changing macroeconomic conditions. This raises an identification problem because age (a_{it}) can be computed as calendar year (t) minus birth year (b_i). Hence, we face the classic problem of distinguishing additive age, year, and cohort effects, where cohort effects are understood to load into the person effects.

In their original paper, Abowd et al. (1999) solved this problem by using actual labor market experience (i.e., the number of years the worker had positive earnings since entering the labor market), which, if some employment histories have gaps, will not be perfectly collinear with year and person dummies. While in some respects this provides a simple fix to the problem, there are two important drawbacks. First, it is not always possible to reconstruct a worker's employment history both because some data sets do not always go far enough back to cover the cohorts of interest and because some data sets report only point-in-time measures of employment (e.g., who was on the payroll in October) rather than a complete history of all employment spells in all years. Second, it is not clear that employment gaps are exogenous, even conditional on a person effect. For example, leaving employment for an entire year could reflect severe health shocks that directly influence earnings ability and confound estimation of relative firm pay.

An alternative approach to dealing with this problem is to impose a linear restriction on the effects of age or time. While the firm effects are invariant to how age and time effects are normalized, different normalizations will yield different values of the person effects and the covariate index $X'_{it}\beta$. Card et al. (2013) allow for separate third-order polynomials in age by education

group along with unrestricted year effects. To obtain identification, they restrict the age profile to be flat at age 40. This is accomplished by omitting the linear age term for each education group and using a cubic polynomial in $a_{it} - 40$. The same restriction is used by Card et al. (2016). While this restriction is unlikely to hold exactly, there is reason to believe it provides a good approximation to the shape of the age-earnings profile.¹¹

Table 3 examines the sensitivity of the results of Card et al. (2016) to four alternate normalizations of the age effects. The first column shows the baseline normalization, which attributes a relatively small fraction of the overall variance of wages to the time-varying individual component of wages. Renormalizing the age profile to be flat at age 50 (col. 2) has little effect on this conclusion, whereas renormalizing the profile to be flat at age 30 leads to a slightly larger variance share for the time-varying component and also implies a relatively strong negative correlation between the person effects and the index $X'_{it}\beta$. Normalizing the age profile to be flat at age 0—which is what is being done by simply omitting the linear term from an uncentered age polynomial—exacerbates this pattern and leads to a decomposition that suggests that the variances of α_i and $X'_{it}\beta$ are both very large and that the two components are strongly negatively correlated.¹² Figure 2 contrasts the implied age profiles for four single year-of-birth cohorts of low-education men from this naive specification, with the implied profiles for the same groups under the baseline normalization. Evidently, the strong negative correlation between the person effects and the covariate index reported in column 4 of table 3 is driven by implausibly large cohort effects, which trend in a way to offset the imposed assumption that the cubic age profile is flat at age 0.

Rather than restricting the age profile to be flat at a point, we can also achieve identification by assuming that the true profile is everywhere nonlinear. Column 5 shows the results of using a linear combination of normal density functions in age (with 5-year bandwidths) to approximate the age profile.¹³ Because each Gaussian component is nonlinear, we do not need restrictions on the parameters to avoid collinearity with cohort and time effects. Nevertheless, using Gaussian basis functions will solve the identification problem only if the true age profile has no linear segments. As shown in column 5, the Gaussian approximation yields results somewhere between our baseline normalization and the specification in column 3: although the

¹¹ For example, as shown in figure 3a–3c of Card and Cardoso (2012), the age profile of wages for Portuguese men tends to be relatively flat after age 40.

¹² Abowd et al. (2003) impose a normalization on the experience profiles in their estimation of an AKM model for the Longitudinal Employer-Household Dynamics data that leads to large variances of the α_i and $X'_{it}\beta$ components and a large negative covariance ($\rho = -0.55$), similar to the pattern in col. 4.

¹³ Letting $n(\cdot)$ denote the standard normal density, we use basis functions of the form $n[(a_{it} - x)/5]$, where $x \in \{20, 25, \dots, 65\}$.

Table 3
Summary of Estimated Abowd, Kramarz, and Margolis (1999) Models for Portuguese Men, Alternative Normalizations of Age Function

	Cubic Age Function Flat				Gaussian Basis Function (5)
	Age 40 (Baseline) (1)	Age 50 (2)	Age 30 (3)	Age 0 (4)	
SD of person effects (across person-year observations)	.42	.41	.46	.93	.44
SD of firm effects (across person-year observations)	.25	.25	.25	.25	.25
SD of Xb (across person-year observations)	.07	.10	.12	.74	.08
Correlation of person/firm effects	.17	.16	.17	.14	.17
Correlation of person effects and covariate index	.19	.19	-.32	-.89	-.06
Correlation of firm effects and covariate index	.11	.14	-.03	-.08	.04
Inequality decomposition (percentage of variance of log wage explained):					
Person effects + covariate index	63	63	63	63	63
Person effects	58	54	70	282	62
Covariate index	2	3	4	180	2
Covariate of person effects and covariate index	3	5	-11	-399	-1
Firm effects	20	20	20	20	20
Covariance of firm effects with person effect + covariate index	12	12	12	12	12
Covariance of firm effects with person effects	11	10	13	21	12
Covariance of firm effects with covariate index	1	2	-1	-9	0
Residual	5	5	5	5	5

NOTE.—The sample includes 8,225,752 person-year observations for male workers in the largest connected set of QP in 2005–9. Sample and baseline specifications are the same as in the study by Card et al. (2016). Models include 1,889,366 dummies for individual workers and 216,459 dummies for individual firms, year dummies interacted with education dummies, and function of age interacted with education dummies. The age function in models in cols. 1–4 includes quadratic and cubic terms, with age deviated from 40, 50, 30, and 0 for models in cols. 1–4, respectively. The age function in model in col. 5 is a Gaussian basis function with five equally spaced spline points. All models have the same fit; root mean square error of the model is 0.143, and the adjusted R^2 is 0.934. SD = standard deviation; Xb = fitted covariate index.

estimated variability of the worker, firm, and time-varying components is very close to baseline, the correlation of the person effects and $X'_{it}\beta$ becomes slightly negative. Fortunately, the covariance of the person and firm effects is essentially the same under our baseline normalization and the Gaussian specification, leading us to conclude that most of the statistics of interest in this literature found under an age 40 normalization are robust to alternate identifying assumptions.

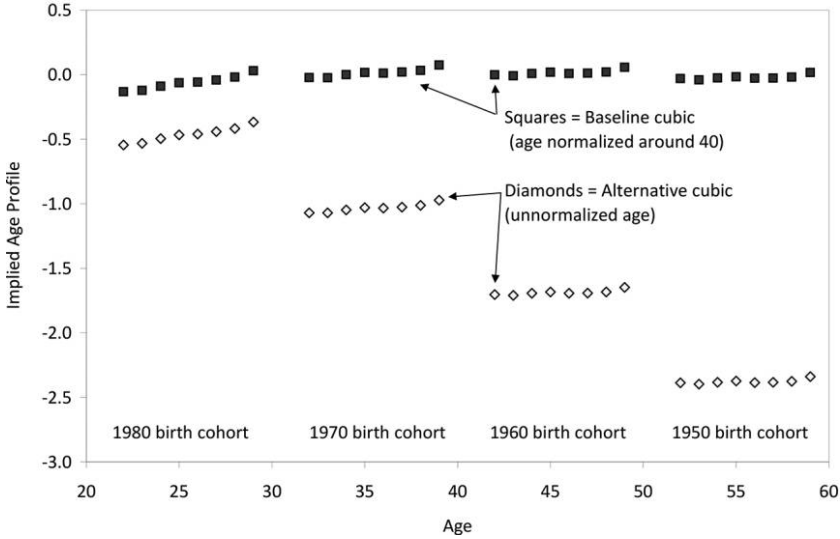


FIG. 2.—Implied age profiles from Abowd, Kramarz, and Margolis (1999) models with alternative normalizations of the age profile (men with primary education only). A color version of this figure is available online.

To summarize: in comparing results from different applications of the AKM framework researchers should pay close attention to the choice of normalization. The values of the person effects (i.e., the α_i 's) and the time-varying controls (i.e., $X_{it}\beta$) are not separately identified when X_{it} includes both year effects and a linear age term. The choice of normalization has no effect on estimates of $\text{var}(\psi_{J(i,t)})$, $\text{var}(\alpha_i + X_{it}\beta)$, or the covariance term $\text{cov}(\psi_{J(i,t)}, \alpha_i + X_{it}\beta)$, but as shown in table 3, it will affect the estimated covariance of the person and firm effects and the relative magnitudes of $\text{var}(\alpha_i)$ and $\text{var}(\psi_{J(i,t)})$.

C. Worker-Firm Sorting and Limited Mobility Bias

In their original study, Abowd et al. (1999) reported a negative correlation between the estimated worker and firm effects, suggesting that sorting of workers to different firms tended to reduce rather than increase overall wage inequality. Subsequent research, however, has typically found positive correlations. For example, Abowd et al. (2003) report a correlation of 0.08 for US workers, while Card et al. (2013) report a correlation of 0.23 for male German workers in the 2000s. As discussed by Abowd et al. (2004) and Andrews et al. (2008), these correlations are biased down in finite samples with the size of the bias depending inversely on the degree of worker mobility among firms. Maré and Hyslop (2006) and Andrews et al. (2012) show convincingly that this limited-mobility bias can be substantial. In sam-

pling experiments they find that the correlation of the estimated effects becomes more negative when the AKM model is estimated on smaller subsets of the available data. While Andrews et al. (2008) and Gaure (2014) provide approaches to correcting for this downward bias in the correlation (and the upward biases in the estimated variances of person and firm effects), their procedures require a complete specification of the covariance structure of the time-varying errors, which makes such corrections highly model dependent.¹⁴ The development of corrections that are robust to unmodeled dependence and heteroscedasticity is an important priority for future research.

D. Exogenous Mobility

Abowd et al.'s (1999) additive worker and firm effect specification is simple and tractable. Nevertheless, it has been widely criticized because OLS estimates of worker and firm effects will be biased unless worker mobility is uncorrelated with the time-varying residual components of wages. In an attempt to provide some transparent evidence on this issue, Card et al. (2013) develop a simple event study analysis of the wage changes experienced by workers moving between different groups of firms. Rather than rely on a model-based grouping, Card et al. (2013) define firm groups on the basis of the average pay of coworkers. If the AKM model is correct and firms offer proportional wage premiums for all their employees, then workers who move to firms with more highly paid coworkers will on average experience pay raises, while those who move in the opposite direction will experience pay cuts. Moreover, the gains and losses for movers in opposite directions between any two groups of firms will be symmetric. In contrast, models of mobility linked to the worker- and firm-specific match component of wages (e.g., Eeckhout and Kircher 2011) imply that movers will tend to experience positive wage gains regardless of the direction of their move, violating the symmetry prediction.

Figures 3 and 4 present the results of this analysis using data for male and female workers in Portugal, taken from Card et al. (2016). The samples are restricted to workers who switch establishments and have at least 2 years of tenure at both the origin and destination firm. Firms are grouped into coworker pay quartiles (using data on male and female coworkers). For clarity, only the wage profiles of workers who move from jobs in quartile 1 and quartile 4 are shown in the figures. The wage profiles exhibit clear steplike patterns: when workers move to higher-paying establishments, their wages

¹⁴ For example, Andrews et al. (2008) compute bias corrections in a linked sample of German workers and establishments under the assumption that the transitory errors in wages are homoscedastic and serially uncorrelated. They find that the corrections have little effect on the estimated correlation between worker and firm effects. However, subsequent results by Andrews et al. (2012) show large biases in the estimated correlation when the AKM model is estimated on subsamples as large as 30% of the data.

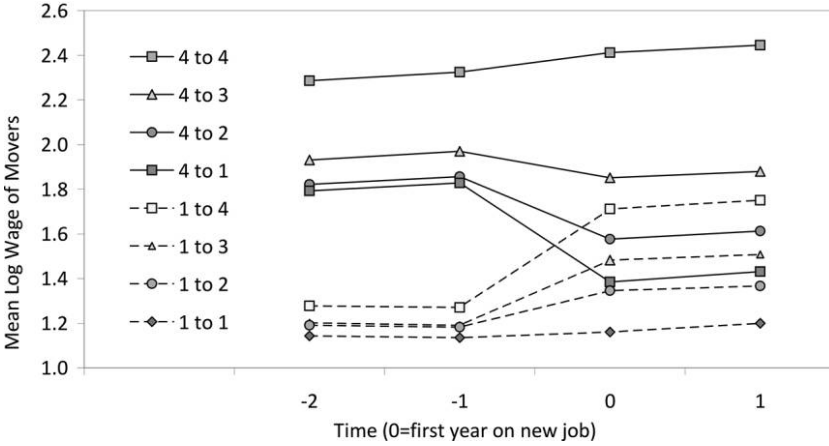


FIG. 3.—Mean log wages of Portuguese male job changers classified by quartile of coworker wages at origin and destination. The figure shows mean wages of male workers at mixed-gender firms who changed jobs in 2004–7 and held the preceding job for 2 years or more and the new job for 2 years or more. Jobs are classified into quartiles based on mean log wage of coworkers of both genders. Source: Card et al. (2016, fig. I). A color version of this figure is available online.

rise; when they move to lower-paying establishments, their wages fall. For example, males who start at a firm in the lowest quartile group and move to a firm in the top quartile have average wage gains of 39 log points, while those who move in the opposite direction have average wage losses of 43 log points. The gains and losses for other matched pairs of moves are also roughly symmetric, while the wage changes for people who stay in the same coworker pay group are close to 0.

Another important feature of the wage profiles in figures 3 and 4 is that wages of the various groups are all relatively stable in the years before and after a job move. Workers who are about to experience a major wage loss by moving to a firm in a lower coworker pay group show no obvious trend in wages beforehand. Similarly, workers who are about to experience a major wage gain by moving to a firm in a higher pay group show no evidence of a pretrend. By contrast, if worker mobility were driven by gradual employer learning, we would expect wage changes to precede moves between firm quality groups over the time horizons examined (Lange 2007).

Card et al. (2016) also present simple tests of the symmetry restrictions imposed by the AKM specification, using regression-adjusted wage changes of males and females moving between firms in the four coworker pay groups. Comparisons of upward and downward movers are displayed visually in figure 5 and show that the matched pairs of adjusted wage changes are roughly scattered along a line with slope of -1 , consistent with the symmetry restriction.

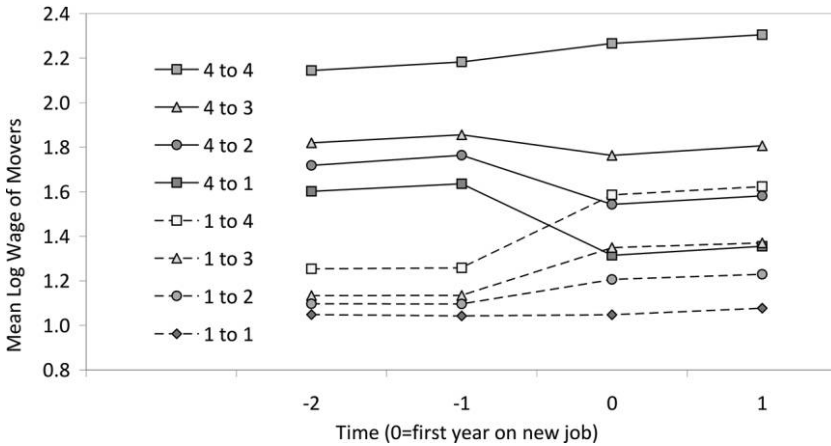


FIG. 4.—Mean wages of Portuguese female job changers classified by quartile of coworker wages at origin and destination. The figure shows mean wages of female workers at mixed-gender firms who changed jobs in 2004–7 and held the preceding job for 2 years or more and the new job for 2 years or more. Jobs are classified into quartiles based on mean log wage of coworkers of both genders. Source: Card et al. (2016, fig. II). A color version of this figure is available online.

Similar figures can be constructed using firm groupings based on the estimated pay effects obtained from an AKM model. As shown in Card et al. (2013, their fig. VII), applying this approach to data for German males yields the same conclusions as an analysis based on coworker pay groups. Macis and Schivardi (2016) report such diagnostics using social security earnings data for Italian workers and confirm that wage profiles of movers exhibit the same steplike patterns found in Germany and Portugal.

E. Additive Separability

Another concern with the AKM model is that it presumes common proportional firm wage effects for all workers. One way to evaluate the empirical plausibility of the additive AKM specification is to examine the pattern of mean residuals for different groups of workers and firms. Figures 6 and 7, taken from Card et al. (2016), show the mean residuals for 100 cells on the basis of deciles of the estimated worker effects and deciles of the estimated firm effects. If the additive model is correct, the residuals should have mean 0 for matches composed of any grouping of worker and firm effects, while if the firm effects vary systematically with worker skill, we expect departures from 0. Reassuringly, the mean residuals are all relatively close to 0. In particular, there is no evidence that the most able workers (in the 10th decile of the distribution of estimated person effects) earn higher premiums at the highest-paying firms (in the 10th decile of the distribution of estimated firm

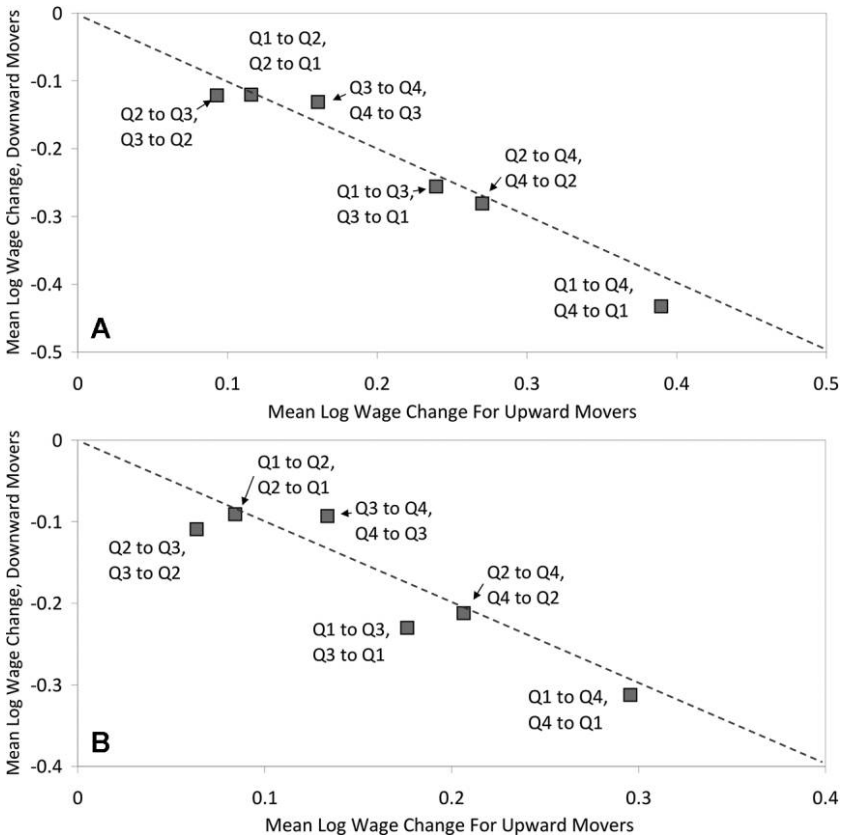


FIG. 5.—A, Test for symmetry of regression-adjusted wage changes of Portuguese male movers across coworker wage quartiles. The figure plots regression-adjusted mean wage changes over a 4-year interval for job changers who move across the coworker wage quartile groups indicated. The dashed line represents symmetric changes for upward and downward movers. Source: Card et al. (2016, fig. B3). B, Test for symmetry of regression-adjusted wage changes of Portuguese female movers across coworker wage quartiles. The figure plots regression-adjusted mean wage changes over a 4-year interval for job changers who move across the coworker wage quartile groups indicated. The dashed line represents symmetric changes for upward and downward movers. Source: Card et al. (2016, fig. B4). A color version of this figure is available online.

effects). The largest mean residuals are for the lowest-ability workers in the lowest paying firms, an effect that may reflect the impact of the minimum wage in Portugal. Residual plots for workers and firms in Germany (reported by Card et al. [2013]) and in Italy (reported by Macis and Schivardi [2016]) also show no evidence of systematic departures from the predictions of a simple AKM-style model.

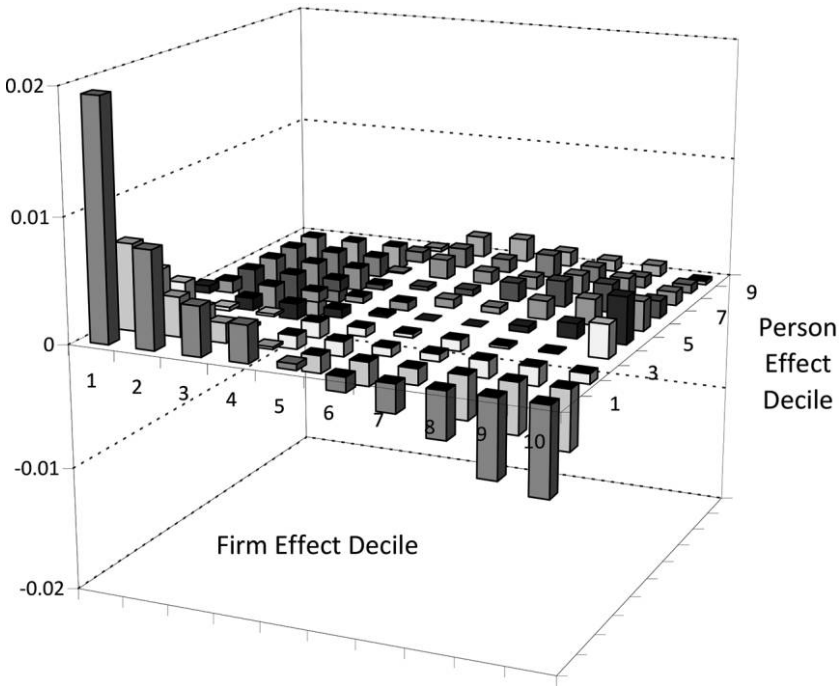


FIG. 6.—Mean residuals by person/firm deciles for Portuguese male workers. The figure shows mean residuals from an estimated Abowd, Kramarz, and Margolis (1999) model with cells defined by decile of estimated firm effects interacted with decile of estimated person effect. Source: Card et al. (2016, fig. B5). A color version of this figure is available online.

A different approach to assessing the additive separability assumption comes from Bonhomme, Lamadon, and Manresa (2015), who estimate a worker-firm model with discrete heterogeneity where each pairing of worker and firm type is allowed a different wage effect. Their results indicate that an additive model provides a very good approximation to Swedish employer-employee data; allowing interactions between worker and firm type yields a trivial (0.8%) increase in explained wage variance. Lochner and Schulz (2016) reach a similar conclusion using information on relative wage rankings inside the firm and firm value added to infer worker and firm types. They find, using German data, that additive separability provides a good approximation to the wage structure, except for the lowest-skilled workers.

Although these results suggest that firm effects are, on average, similar for different types of workers, there is of course scope for differences to emerge in selected subpopulations. For example, Goldschmidt and Schmieler (2015) find in large German firms that food, cleaning, security, and logistics (FCSL)

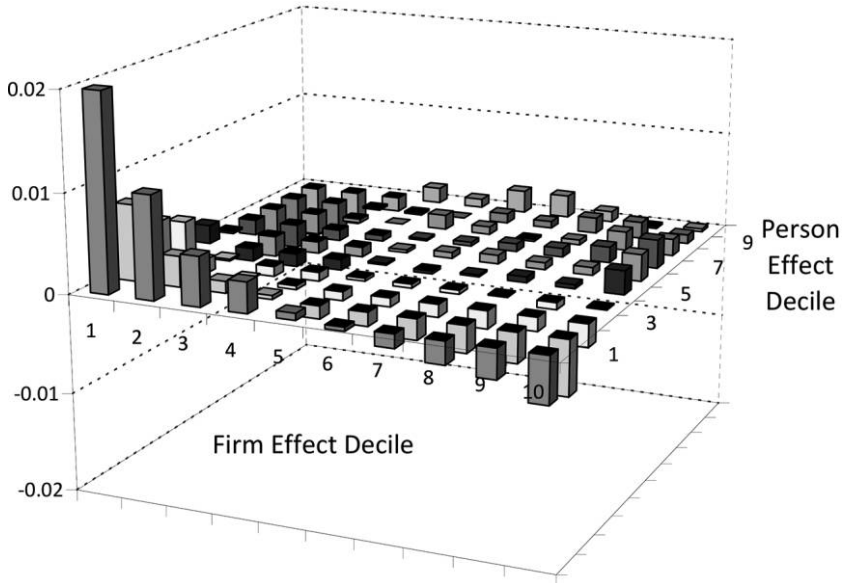


FIG. 7.—Mean residuals by person/firm deciles for Portuguese female workers. The figure shows mean residuals from an estimated Abowd, Kramarz, and Margolis (1999) model with cells defined by decile of estimated firm effects interacted with decile of estimated person effect. Source: Card et al. (2016, fig. B6). A color version of this figure is available online.

workers exhibit different wage fixed effects than other occupations. Specifically, the firm wage effects of FCSL workers are attenuated relative to non-FCSL workers. Likewise, Card et al. (2016) find that Portuguese women exhibit slightly attenuated firm effects relative to men, which they argue reflects gender differences in bargaining behavior.

IV. Reconciling Rent-Sharing Estimates with Results from Studies of Firm Switching

In their original study, Abowd et al. (1999) showed that the estimated firm-specific wage premiums were positively correlated with measures of firm profitability, including value added per worker and sales per worker. A number of more recent studies have also confirmed that there is a positive link between firm-specific pay policies and productivity (e.g., Cahuc, Postel-Vinay, and Robin 2006; Bagger, Christensen, and Mortensen 2014).

To further bridge the gap between the rent-sharing literature and the firm-wage effects literatures, we conducted a simple exercise using data on male workers in Portugal observed in the QP between 2005 and 2009 (i.e., the same data used in panel A of table 2). The AKM model posits that the log of the

wage of a given worker in a given year can be decomposed into the sum of a person effect, a firm or establishment effect, a time-varying index of person characteristics, and a residual that is orthogonal to the firm and person effects. It follows that the rent-sharing elasticity obtained from a regression of wages on a time-invariant measure of rents at the current employer (γ_w) can be decomposed into the sum of three components reflecting the regression on firm-specific rents of the estimated worker effects (γ_α), the estimated firm effects (γ_ψ), and the time-varying covariate index ($\gamma_{X\beta}$):

$$\gamma_w = \gamma_\alpha + \gamma_\psi + \gamma_{X\beta}.$$

The regression coefficients γ_α and $\gamma_{X\beta}$ represent sorting effects. To the extent that firms with larger measured rents hire older workers or workers with greater permanent skills, γ_α and/or $\gamma_{X\beta}$ will be positive. The coefficient γ_ψ , on the other hand, is arguably a clean measure of the rent-sharing elasticity, since $\psi_{J(i,t)}$ represents a firm-specific wage premium that is paid on top of any reward for individual-specific skills.

To implement this idea, we use the estimated AKM parameters from Card et al. (2016), which were estimated on a sample that includes virtually all the observations used for the cross-sectional models in panel A of table 2.¹⁵ The results are presented in panel A of table 4. Row 1 of the table reports estimated rent-sharing elasticities using the log hourly wage of each worker as a dependent variable. As in table 2, we report three specifications corresponding to models with only simple human capital controls (col. 1), controls for major industry and city (col. 2), and controls for detailed industry and location (col. 3). The estimated rent-sharing elasticities in row 1 are qualitatively similar to the estimates in row 1 of table 2 but differ slightly because the AKM model estimates are not available for all workers/firms. Rows 2–4 show how the overall rent-sharing elasticities in row 1 can be decomposed into a worker quality effect (row 2), a firm wage premium effect (row 3), and an experience-related sorting effect (row 4), which is close to 0.

A key conclusion from these estimates is that rent-sharing elasticities estimated from a cross-sectional specification incorporate a sizable worker quality bias. In each column of table 4, roughly 40% of the overall wage elasticity in row 1 is due to the correlation of worker quality (measured by the person effect component of wages) with firm-specific quality. Adjusting for worker quality, the estimates in row 3 point to a rent-sharing elasticity in the range of 0.10–0.15, large enough to create a Lester range of wage variation of 16–24 log points associated with the differences between firms at the 90th and 10th percentiles of log value added per worker.

¹⁵ The sample used by Card et al. (2016) is slightly different than the sample of firms with financial data that we use in this paper, so the adding-up constraint does not have to hold exactly. However, in all cases it holds approximately.

Table 4
Relationship between Components of Wages and Mean Log Value Added per Worker

	Basic Specification (1)	Basic + Major Industry/City (2)	Basic + Detailed Industry/City (3)
A. Combined sample ($n = 2,252,436$ person-year observations at 41,120 firms):			
Log hourly wage	.250 (.018)	.222 (.016)	.187 (.012)
Estimated person effect	.107 (.010)	.093 (.009)	.074 (.006)
Estimated firm effect	.137 (.011)	.123 (.009)	.107 (.008)
Estimated covariate index	.001 (.000)	.001 (.000)	.001 (.000)
B. Less educated workers ($n = 1,674,676$ person-year observations at 36,179 firms):			
Log hourly wage	.239 (.017)	.211 (.016)	.181 (.011)
Estimated person effect	.089 (.009)	.072 (.009)	.069 (.005)
Estimated firm effect	.144 (.015)	.133 (.013)	.107 (.008)
Estimated covariate index	.000 (.000)	.000 (.000)	.000 (.000)
C. More educated workers ($n = 577,760$ person-year observations at 17,615 firms):			
Log hourly wage	.275 (.024)	.247 (.020)	.196 (.017)
Estimated person effect	.137 (.016)	.130 (.013)	.094 (.009)
Estimated firm effect	.131 (.012)	.113 (.009)	.099 (.010)
Estimated covariate index	-.001 (.000)	-.001 (.000)	-.001 (.000)

NOTES.—Entries are coefficients of mean log value added per worker (at current firm) in regression models with dependent variables listed in the row headings. Standard errors are clustered by firm (in parentheses). The sample in panel B includes males with less than completed secondary education at firms in the connected set for less educated workers. The sample in panel C includes males with a high school education or more at firms in the connected set for more educated workers. The sample in panel A includes males in either the panel B or the panel C sample. All models control for cubic in experience and unrestricted education \times year dummies. Models in col. 2 also control for 20 major industries and two major cities (Lisbon and Porto). Models in col. 3 also control for 202 detailed industry dummies and 29 Nomenclature of Territorial Units for Statistics region 3 location dummies.

While the AKM approach reduces the estimated rent-sharing elasticities substantially, the estimates in row 3 of table 4 are still substantially larger than the within-job elasticities reported in panel B of table 2. There are several possible explanations for this gap. One is that the within-job estimates are biased downward by measurement errors, which comprise a potentially

large share of the variance in relatively short-horizon changes in rents. A related explanation, emphasized by Guiso et al. (2005), is that wages tend to adjust less to purely transitory fluctuations than to persistent changes in productivity.¹⁶ To the extent that industry-wide productivity shifts are more persistent than firm-specific within-industry shifts, this explanation can also account for the pattern of smaller elasticities when more detailed industry controls are added to a rent-sharing model.

A third explanation is that some share of the firm-specific wage premium paid by more productive firms is a compensating differential for extra work effort or less desirable working conditions (Rosen 1986; Hwang, Mortensen, and Reed 1998). How much of the gap between firm-stayer designs and worker-switching designs such motives can explain remains an open question (Bonhomme and Jolivet 2009; Sorkin 2015; Lavetti and Schmutte 2016; Taber and Vejlín 2016). In a recent field experiment, Mas and Pallais (2016) find that workers are not willing to pay much for scheduling flexibility but are willing to pay to avoid working evenings and weekends. In line with this evidence, Card et al. (2016, their table B6) examine the relationship between average hours of work and the estimated pay premiums offered by different firms in Portugal and find no evidence of compensating differentials for long hours. Nevertheless, we cannot rule out some role for compensating differentials, which would imply that the estimates in row 3 of table 4 may overstate the true rent-sharing elasticity. What is clear is that jobs with higher wage premiums last significantly longer (Card, Heining, and Kline 2012, their table 8), which we take as providing relatively strong evidence that the AKM firm effect estimates capture a significant rent component.

A. Differential Rent Sharing

We can use the AKM framework to examine another interesting question: to what extent do different groups of workers receive larger or smaller shares of the rents at different firms? To do this, we fit separate AKM models for less educated men (with less than a high school education) and more educated men (with a high school education or more) to our Portuguese wage sample. We then reestimated the same rent-sharing specifications reported in panel A of table 4 separately for the two groups. The results are reported in panels B and C of table 4.

The estimates reveal several interesting patterns. Most importantly, although the correlation between wages and value added per worker is a little higher for the more educated men, virtually all of this gap is due to a stronger correlation between the worker quality component of wages and value

¹⁶ Cardoso and Portela (2009) find evidence for this pattern using Portuguese worker-firm data derived from the QP.

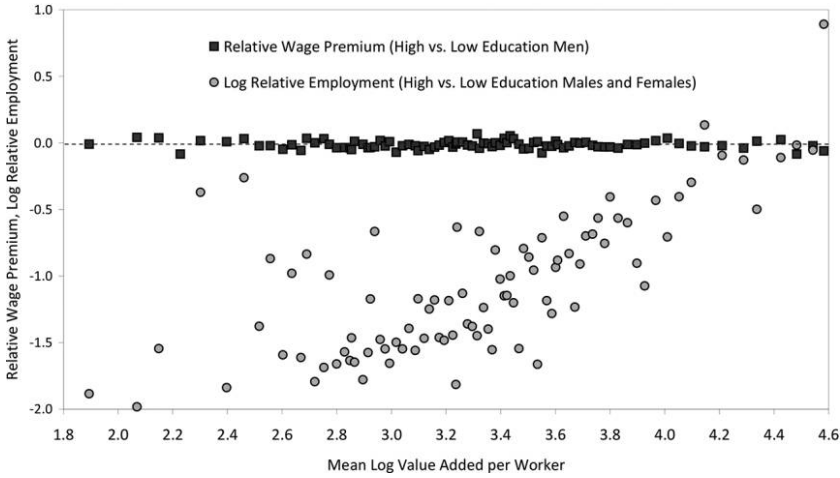


FIG. 8.—Relative wage premium and relative employment of high- versus low-education workers. Firms are divided into 100 cells on the basis of mean log value added per worker in 2005–9, with equal numbers of person-year observations per cell. A color version of this figure is available online.

added. The correlations with the firm-specific pay premiums are very similar for the two education groups. Thus, we see no evidence of differential rent sharing.

This finding is illustrated in figure 8, which shows a binned scatterplot of mean log value added per worker at different firms (on the horizontal axis) versus the relative wage premium for high-educated versus low-educated men at these firms. We also superimpose a bin-scatter of the relative share of higher-education workers at different firms (including both men and women in the employment counts for the two education groups). The relative wage premium is virtually flat, consistent with the regression coefficients in rows 7 and 11 of table 4, which show nearly the same effect of value added per worker on the wage premiums for the two education groups. In contrast, the relative share of highly educated workers is increasing with value added per worker, a pattern we interpret as largely driven by the labor quality component in value added per worker.¹⁷

¹⁷ As discussed in Sec. I, value added per worker will in general depend on the quality of the workforce. For example, if $VA_j = P_j T_j [(1 - \theta)L_j + \theta H_j]$, where L_j and H_j are the numbers of low- and high-skilled workers employed at the firm, and $N_j = L_j + H_j$, then $\ln(TFP_j/N_j) = \ln(TFP_j) + \ln(q_j)$, where $q_j = [(1 - \theta)L_j + \theta H_j]/N_j$ is a measure of the quality of the workforce at firm j . The expected slope of a regression of the log of the relative share of highly educated workers on the log of value added per worker is therefore positive, even if there is no correlation between TFP and the share of highly educated workers.

V. Imperfectly Competitive Labor Markets and Inequality

With this background in mind, we now turn to the task of developing a simple modeling framework that is useful for organizing and interpreting the empirical literature on firm-specific productivity and wage dispersion. In contrast to much of the rent-sharing literature, we assume that wages are set by employers to maximize profits, subject to constraints on the relationship between wages and the supply of labor. Rather than build a model of the supply side based on search frictions, we follow the industrial organization literature by working with a static differentiated products model that focuses on heterogeneity across workers in their valuation of jobs at different employers. This differentiation endows firms with some power to set wages as in classic monopsony models.¹⁸

While empirical work on monopsony has experienced something of a renaissance (for a review, see Manning 2011), to our knowledge there has been little attempt to use these models to reconcile facts in the literature on matched employer-employee data. We show that static monopsony models can generate empirically plausible connections between firm productivity and wages, observationally equivalent to simple rent-sharing models such as equation (1). They also, under reasonable assumptions, generate the prediction that wages are additively separable in worker and firm heterogeneity, at least within broad skill groups.

A limitation of our framework relative to modern wage-posting models is that we assume that all between-firm heterogeneity arises from heterogeneity in TFP or differences in the elasticity of labor supply to the firm. While this allows us to focus on the links between dispersion in productivity and wages, it is important to remember that firms may also exhibit dispersion in wage policies for reasons having nothing to do with their production technology. Indeed, in the simplest version of Burdett and Mortensen's (1998) model, firms are homogenous, and the identity of high-wage and low-wage firms is arbitrary.¹⁹

¹⁸ In this respect, our approach is akin to the classic Albrecht-Axell (1984) model of wage posting with leisure heterogeneity. However, because we allow for continuous heterogeneity in worker preferences, firms are not indifferent between wage strategies and will mark wages down below marginal product, according to the usual monopsonistic pricing rule. Our assumption that firms are ignorant about worker reservation values lies in contrast to the model of Postel-Vinay and Robin (2002), who assume that firms observe a worker's outside option and offer wages that make them indifferent about accepting jobs.

¹⁹ We have also ignored efficiency wage explanations for firm wage premia, which can emerge, e.g., as a result of monitoring problems. See Akerlof and Yellen (1986) and Katz (1986) for reviews and Piyapromdee (2013) for an attempt to combine efficiency wage mechanisms with wage-posting models.

A. Market Structure

There are J firms and two types of workers: lower skilled (L) and higher skilled (H). Each firm $j \in \{1, \dots, J\}$ posts a pair (w_{Lj}, w_{Hj}) of skill-specific wages that all workers costlessly observe. Hence, in contrast to search models, workers are fully informed about job opportunities. As in many search models, however, we assume that firms will hire any worker (of appropriate quality) who is willing to accept a job at the posted wage.

Firms exhibit differentiated work environments over which workers have heterogeneous preferences. For worker i in skill group $S \in \{L, H\}$, the indirect utility of working at firm j is

$$u_{iSj} = \beta_S \ln(w_{Sj} - b_S) + a_{Sj} + \epsilon_{iSj},$$

where b_S is a skill group-specific reference wage level (e.g., arising from wages paid in an outside competitive sector), a_{Sj} is a firm-specific amenity common to all workers in group S , and ϵ_{iSj} captures idiosyncratic preferences for working at firm j , arising, for example, from nonpecuniary match factors such as distance to work or interactions with coworkers and supervisors. We assume that the $\{\epsilon_{iSj}\}$ are independent draws from a type I extreme value distribution.²⁰ Given posted wages, workers are free to work at any firm they wish. Hence, by standard arguments (McFadden 1973), workers have logit choice probabilities of the form

$$p_{Sj} \equiv P \left(\arg \max_{k \in \{1, \dots, J\}} \{u_{iSk}\} = j \right) = \frac{\exp(\beta_S (\ln(w_{Sj} - b_S) + a_{Sj}))}{\sum_{k=1}^J \exp(\beta_S (\ln(w_{Sk} - b_S) + a_{Sk}))}.$$

To simplify the analysis and abstract from strategic interactions in wage setting, we assume that the number of firms J is very large, in which case the logit probabilities are closely approximated by exponential probabilities

$$p_{Sj} \approx \lambda_S \exp(\beta_S \ln(w_{Sj} - b_S) + a_{Sj}),$$

where (λ_H, λ_L) are constants common to all firms in the market. Thus, for large J , the approximate firm-specific supply functions are

$$\ln L_j(w_{Lj}) = \ln(\mathcal{L}\lambda_L) + \beta_L \ln(w_{Lj} - b_L) + a_{Lj}, \tag{5}$$

$$\ln H_j(w_{Hj}) = \ln(\mathcal{H}\lambda_H) + \beta_H \ln(w_{Hj} - b_H) + a_{Hj}, \tag{6}$$

²⁰ A potentially important distinction across worker subgroups is the relative magnitude of their variation in idiosyncratic preferences. For example, assume that $u_{iSj} = \beta_S^0 \ln(w_{Sj} - b_S) + a_{Sj} + \tau_S \epsilon_{iSj}$, where τ_S is larger for groups with a wider variation in idiosyncratic preferences and ϵ_{iSj} is again extreme value distributed. This gives rise to the same choice probabilities as preferences $u_{iSj} = (\beta_S^0/\tau_S) \ln(w_{Sj} - b_S) + (1/\tau_S)a_{Sj} + \epsilon_{iSj}$.

where \mathcal{L} and \mathcal{H} give the total numbers of lower-skilled and higher-skilled workers in the market.²¹ Note that as $\beta_L, \beta_H \rightarrow \infty$ these supply functions become perfectly elastic, and we approach a competitive labor market with exogenous wages b_L and b_H .

B. Firm Optimization

Firms have production functions of the form

$$Y_j = T_{ij}f(L_j, H_j), \tag{7}$$

where T_i is a firm-specific productivity shifter. We assume that $f(\dots)$ is twice differentiable and exhibits constant returns to scale with respect to L_j and H_j . For simplicity, we also ignore capital and intermediate inputs.²²

The firm’s problem is to post a pair of skill-specific wages that minimize the cost of labor services given knowledge of the supply functions (5) and (6). Firms cannot observe workers’ preference shocks $\{\epsilon_{isj}\}$, which prevents them from perfectly price discriminating against workers according to their idiosyncratic reservation values. The firm’s optimal wage choices solve the cost minimization problem

$$\min_{w_L, w_H} w_L L_j(w_L) + w_H H_j(w_H) \text{ such that } T_{ij}f(L_j(w_L), H_j(w_H)) \geq Y.$$

The associated first-order conditions can be written as

$$w_L \frac{1 + e_{Lj}}{e_{Lj}} = T_{ij}f_L \mu_j, \tag{8}$$

$$w_H \frac{1 + e_{Hj}}{e_{Hj}} = T_{ij}f_H \mu_j, \tag{9}$$

where e_{Lj} and e_{Hj} represent the elasticities of supply of L and H workers at the optimal choice of wages and μ_j represents the marginal cost of produc-

²¹ Berry and Pakes (2007) contrast demand models, where consumers have idiosyncratic preferences for specific products, with what they term the “pure characteristics” model, where consumers care about only a finite set of product characteristics. In the latter case, as the number of products grows large, the demand elasticity tends to infinity, a phenomenon discussed in the labor market setting by Boal and Ransom (1997). We suspect that the pure characteristics model is less applicable to the worker’s choice of employer because of the many nonpecuniary aspects of work that can give rise to match effects. For example, no two employers have exactly the same location and workplace culture. However, which model works better empirically is clearly an important question for future research.

²² This specification is appropriate if the user cost of capital and the prices of intermediate inputs are fixed and the firm’s output is a Cobb-Douglas function of these factors and the labor aggregate $T_{ij}f(L_j, H_j)$. In this case, capital and intermediate inputs will adjust proportionally to $T_{ij}f(L_j, H_j)$.

tion, which the firm will equate to marginal revenue at an optimal choice for Y . Thus, the terms $T_{ij}f_{Lj}\mu_j$ and $T_{ij}f_{Hj}\mu_j$ on the right-hand sides of equations (8) and (9) represent the marginal revenue products of the two types of labor, while the terms on the left-hand sides represent their marginal factor costs. Using equations (5) and (6), the elasticities of supply are

$$e_{Lj} = \frac{\beta_L \omega_{Lj}}{\omega_{Lj} - b_L},$$

$$e_{Hj} = \frac{\beta_H \omega_{Hj}}{\omega_{Hj} - b_H}.$$

Note that for both groups, labor supply to the firm becomes infinitely elastic as wages approach the reference wage level b_s . Using these expressions, the firm's first-order conditions can be rewritten as

$$\omega_{Lj} = \frac{1}{1 + \beta_L} b_L + \frac{\beta_L}{1 + \beta_L} T_{ij}f_{Lj}\mu_j, \tag{10}$$

$$\omega_{Hj} = \frac{1}{1 + \beta_H} b_H + \frac{\beta_H}{1 + \beta_H} T_{ij}f_{Hj}\mu_j. \tag{11}$$

The firm's optimal wage choice for skill group S is a weighted average of the reference wage b_s and the group's marginal revenue product. As $\beta_s \rightarrow \infty$, $e_{sj} \rightarrow \infty$, and the firm pays the reference wage b_s . In the case where the reference wage is 0, the labor supply function for group S has a constant elasticity β_s , and the first-order condition sets the wage equal to a constant fraction $\beta_s/(1 + \beta_s)$ of the marginal revenue product.

Note that firms post wages with knowledge of the shape of the skill-specific supply schedules but not the identities of the workers who comprise them. The last worker hired is indifferent about taking the job, but the other employees strictly prefer their job to outside alternatives. These inframarginal workers capture rents by means of an information asymmetry: they hide from their employer the fact that they would be willing to work for a lower wage.

We now derive results under various specifications of the production function and the firm's marginal revenue function. On the technology side, we start with a simple baseline case where $f(\dots)$ is linear in L_j and H_j . This corresponds to a standard efficiency units model of the labor market in which lower- and higher-skilled workers are perfect substitutes. We then consider the more general case where $f(\dots)$ is a constant elasticity of substitution (CES) production function. On the revenue side, we initially assume that the firm faces a fixed output price. We then consider the more general case where the firm faces a constant elasticity product demand function.

C. Baseline Case: Linear Production Function and Fixed Output Price

To develop intuition, we begin with the simplest possible example, where the firm faces a fixed output price P_j^o and has a linear production function

$$Y_j = T_j N_j = T_j((1 - \theta)L_j + \theta H_j).$$

Here N_j represents the efficiency units of labor at the firm and the parameter $\theta \in (0.5, 1)$, which we assume is common to all firms, governs the relative productivity of the two types of labor. Under this specification of technology and market structure, the first-order conditions (10) and (11) evaluate to

$$\begin{aligned} \omega_{L_j} &= \frac{1}{1 + \beta_L} b_L + \frac{\beta_L}{1 + \beta_L} T_j P_j^o (1 - \theta), \\ \omega_{H_j} &= \frac{1}{1 + \beta_H} b_H + \frac{\beta_H}{1 + \beta_H} T_j P_j^o \theta. \end{aligned}$$

The determination of the optimal wage in the simplified situation where there is only one skill group is illustrated in figure 9. The firm faces an upward-sloping inverse labor supply function of the form $\omega = b + N^{1/\beta}$. The associated marginal factor cost is $MFC = b + [(1 + \beta)/\beta]N^{1/\beta}$. The firm equates MFC with marginal revenue product (MRP), leading to an equilibrium wage $\omega = [1/(1 + \beta)]b + [\beta/(1 + \beta)]MRP$. As shown in the figure, if the firm’s marginal revenue product increases, both employment and wages will increase at the firm. In contrast to traditional rent-sharing

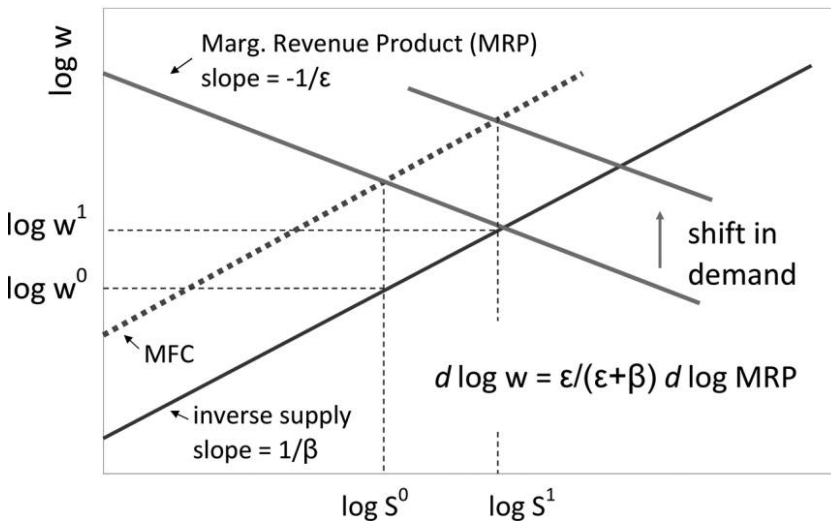


FIG. 9.—Effect of total factor productivity shock (single skill group). MFC = marginal factor cost. A color version of this figure is available online.

models, however, this positive relationship between wages and productivity does not stem from wage bargaining. Firms unilaterally post profit-maximizing wages that leave the marginal worker with no surplus on the job. The firm shares rents with inframarginal workers only because it lacks the information necessary to price discriminate on the basis of reservation wages.

To understand the implications of this model for the relative wage structure, suppose that the reference wages of the two skill groups are proportional to their relative productivities, so that

$$b_L = (1 - \theta)b, \quad b_H = \theta b.$$

This restriction is natural if one views b_s as an outside wage that can be earned in a fully competitive sector where wages equal marginal products. Now the first-order conditions can be rewritten as

$$\ln w_{Lj} = \ln \frac{(1 - \theta)b}{1 + \beta_L} + \ln(1 + \beta_L R_j), \tag{12}$$

$$\ln w_{Hj} = \ln \frac{\theta b}{1 + \beta_H} + \ln(1 + \beta_H R_j), \tag{13}$$

where $R_j \equiv T_j P_j^o / b$ gives the proportional gap in marginal labor productivity at firm j relative to the competitive sector. Wages of both skill groups contain a rent-sharing component that depends on R_j and the skill group-specific supply parameter β_s .

Note that under the linear technology assumption, value added per standardized unit of labor is $v_j \equiv P_j^o Y_j / N_j = P_j^o T_j$, so $R_j = v_j / b$ is the ratio of value added per standardized unit of labor to the outside wage for a worker with 1 efficiency unit of labor. Equations (12) and (13) therefore imply that the elasticity of wages of skill group S with respect to value added per worker is

$$\xi_{sj} \equiv \frac{\partial \ln w_{sj}}{\partial \ln v_j} = \frac{\beta_s R_j}{1 + \beta_s R_j}.$$

Interestingly, this is the same as the expression for the rent-sharing elasticity (eq. [2]) in a bargaining model where workers are assumed to capture a fixed share of the quasi rents.

The estimated rent-sharing elasticities in table 1 (and in our reanalysis of Portuguese data) indicate that a typical value of this elasticity is around 0.10, which suggests that the average value of $\beta_s R_j$ is also around 0.10.²³ Hence, an

²³ This is an approximation because the model implies that rent-sharing elasticities vary systematically across firms in relationship to productivity. Supposing, for simplicity, that an average elasticity is being reported in empirical studies, the approximation that $(1/J) \sum_j \beta_s R_j \approx (1/J) \sum_j \xi_{sj}$ will be accurate when the distribution of ξ_{sj} is bounded from above by a number far below 1.

average worker earns about 10% higher wages than he or she would earn at the lowest-wage firms in the economy that have $v_j = b$. This estimate of average rents earned per worker is remarkably consistent with the estimates of Card et al. (2016, their table III) obtained by benchmarking wage premiums at firms in the Portuguese economy relative to the premiums paid by the least profitable firms in the country. By contrast, Hornstein et al. (2011) show that a wide class of search models have difficulty generating mean wages more than 5% above the lowest offered wage.

The elasticity of labor supply for skill group S when wages are determined by the first-order conditions (12) and (13) is

$$e_{Sj} = \frac{\beta_L \omega_{Sj}}{\omega_{Sj} - b_S} = \frac{1 + \beta_S R_j}{R_j - 1}.$$

Assuming that $\beta_S R_j = 0.1$, a value of the firm-specific elasticity of supply of around 4 implies that $R_j \approx 1.3$ and $\beta_S \approx 0.08$. While many empirical estimates of the elasticity of supply to the firm are lower than 4 (Manning 2011), we consider this a reasonable near-competitive benchmark because it implies an equilibrium markdown of wages relative to marginal products of only 20%.

A key implication of equations (12) and (13) is that when $\beta_L = \beta_H$, the relative wages of the two skill groups are independent of firm-specific productivity. To simplify the discussion, assume that $\beta_L R_j$ and $\beta_H R_j$ are both relatively small (i.e., on the order of 0.10). In such a case, the Taylor approximation

$$\begin{aligned} \ln w_{Lj} &= \ln \frac{(1 - \theta)b}{1 + \beta_L} + \beta_L R_j, \\ \ln w_{Hj} &= \ln \frac{\theta b}{1 + \beta_H} + \beta_H R_j, \end{aligned}$$

will be highly accurate. This implies that the log wage gap between high- and low-skilled workers at firm j is

$$\ln \frac{w_{Hj}}{w_{Lj}} = \ln \frac{\theta}{1 - \theta} + \ln \frac{1 + \beta_L}{1 + \beta_H} + (\beta_H - \beta_L) R_j. \tag{14}$$

When $\beta_L = \beta_H = \beta$, wages can be written in the form

$$\ln w_{Sj} = \alpha_S + \psi_j, \tag{15}$$

where $\alpha_S \equiv \ln(b/(1 + \beta)) + 1(S = L) \times \ln 1 - \theta + 1(S = H) \times \ln \theta$ is a skill group-specific constant and $\psi_j = \beta R_j = (\beta/b)v_j$ is the firm-specific wage premium paid by firm j . This simple model therefore yields a reduced form specification for individual wages that is consistent with the additively separable formulation proposed by Abowd et al. (1999). Moreover, the firm

effects should be strongly related to value added per worker, something we saw evidence for in table 4.²⁴

When one group has a higher value of the supply parameter β , the log wage gap between workers in different skill groups will be higher at more profitable firms. In this case, the data will be described by an AKM-style model with skill group-specific firm effects. The wage premium for skill group S at firm j will be

$$\psi_j^S = \beta_S R_j.$$

The value of the ratio β_H/β_L can be identified from a projection of the firm effects for workers in group H on the associated firm effects for workers in group L (since $\psi_j^H = (\beta_H/\beta_L)\psi_j^L$). Card et al. (2016, their fig. V) relate gender-specific firm effect estimates to one another and find results consistent with a value of β about 10% larger for males than females.

1. Between-Firm Sorting

Even when $\beta_L = \beta_H$ and the wage gap between workers in the two skill groups is constant at any given firm, the market-wide average wage for each skill group will depend on their relative distribution across firms. In particular, equation (15) implies that the expected log wage for workers in skill group S is

$$E[\ln w_{Sj}] = \alpha_S + \sum_j \psi_j \pi_{Sj},$$

where π_{Sj} is the share of workers in skill group S employed at firm j . Thus, the market-wide wage differential between high- and low-skilled workers depends on their relative productivity, their relative supply elasticities, and the relative shares of the two groups employed at firms with higher or lower wage premiums:

$$E[\ln w_{Hj}] - E[\ln w_{Lj}] = \alpha_H - \alpha_L + \sum_j \psi_j (\pi_{Hj} - \pi_{Lj}). \tag{16}$$

The third term in this expression represents a between-firm sorting component of the average wage gap. Card et al. (2016) show that 15%–20% of the wage differential between men and women in Portugal is explained by the fact that males are more likely to work at firms that pay higher wage premiums to both gender groups. Similarly, Card et al. (2012) show that an important share of the rising return to education in Germany is explained by the increasing likelihood that higher-educated workers are sorted to establishments with higher pay premiums.

²⁴ While table 4 regresses the firm effects on $\ln v_j$, the above analysis suggests that ψ_j should be nearly linear in v_j , with a coefficient of approximately β/b . In practice, Card et al. (2016) find that AKM firm effects are instead approximately linear in $\ln v_j$ above a minimal threshold level.

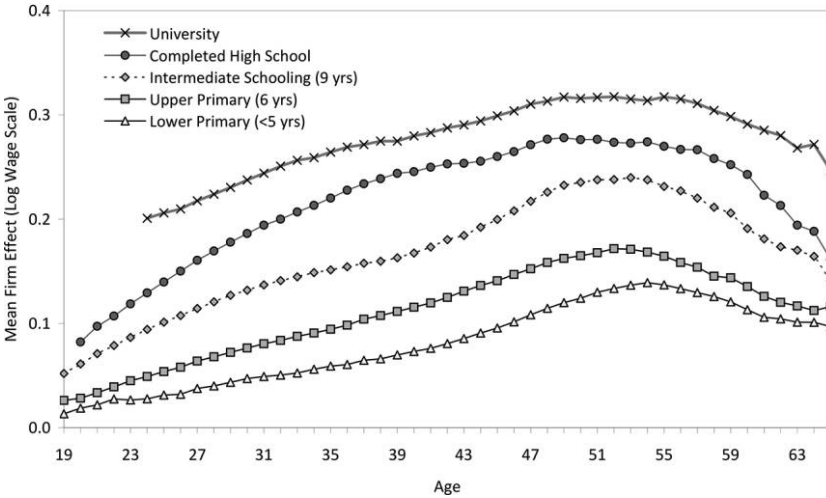


FIG. 10.—Mean firm effects by age and education group for Portuguese males. Firm effects are normalized using the method of Card et al. (2016).

Some simple evidence on the importance of the sorting component for the structure of wages for Portuguese male workers is presented in figure 10. Here, we plot the mean firm effects by age for Portuguese men in five different education groups. We normalize the estimated firm effects using the procedure described by Card et al. (2016), which sets the average firm effect to 0 for firms in (roughly) the bottom 15% of the distribution of log value added per worker. The figure shows two important features. First, within each education group, the mean firm effect associated with the jobs held by workers at different ages is increasing until about age 50 and then slightly decreasing.²⁵ Thus, the life-cycle pattern of between-firm sorting contributes to the well-known shape of the life-cycle wage profile. Second, at all ages more highly educated workers are more likely to work at firms that pay higher wage premiums to all their workers. A significant share of the wage gap between men with different education levels is therefore attributable to differential sorting.

When the supply parameter β varies across groups, the wage decomposition will contain an additional term, reflecting a weighted average across firms of the rent-sharing components of the two skill groups:

²⁵ Topel and Ward (1992) showed that job-to-job mobility was an important component of wage growth for young men in the US labor market. They interpreted their finding as mainly arising from gains in the job-match component of wages rather than as systematic mobility to firms that pay higher wages to all workers.

$$\begin{aligned} E[\ln w_{Hi}] - E[\ln w_{Li}] &= \alpha_H - \alpha_L + \sum_j \psi_j^L (\pi_{Hj} - \pi_{Lj}) + \sum_j (\psi_j^H - \psi_j^L) \pi_{Hj} \\ &= \alpha_H - \alpha_L + \sum_j \psi_j^H (\pi_{Hj} - \pi_{Lj}) + \sum_j (\psi_j^H - \psi_j^L) \pi_{Lj}. \end{aligned}$$

As in a traditional Oaxaca (1973)-style decomposition, these expressions give alternative ways to evaluate the contributions of differences in the distributions of the two skill groups across firms and differences in the return to working at a given firm between the two groups.

D. Downward-Sloping Firm-Specific Product Demand

So far we have assumed that the firm is a price taker in its output market. Suppose now that the firm faces an inverse demand function $P_j = P_j^0 Y_j^{-1/\varepsilon}$, with $\varepsilon > 1$ giving the elasticity of product demand. This yields the marginal revenue function

$$MR_j = \left(\frac{\varepsilon - 1}{\varepsilon} \right) P_j^0 Y_j^{-1/\varepsilon}.$$

In this case, assuming as above that $b_L = (1 - \theta)b$ and $b_H = \theta b$, the first-order conditions (10) and (11) evaluate to

$$w_{Lj} = \frac{b(1 - \theta)}{1 + \beta_L} \left[1 + \beta_L \left(\frac{\varepsilon - 1}{\varepsilon} \right) T_j P_j^0 Y_j^{-1/\varepsilon} \right], \tag{17}$$

$$w_{Hj} = \frac{\theta b}{1 + \beta_H} \left[1 + \beta_H \left(\frac{\varepsilon - 1}{\varepsilon} \right) T_j P_j^0 Y_j^{-1/\varepsilon} \right]. \tag{18}$$

These equations can be simplified by noting that value-added per efficiency unit of labor is

$$v_j \equiv \frac{P_j Y_j}{N_j} = P_j^0 T_j Y_j^{-1/\varepsilon}.$$

Thus, the optimal choices for wages can be written

$$\ln w_{Lj} = \ln \frac{(1 - \theta)b}{1 + \beta_L} + \ln(1 + \beta_L R_j'),$$

$$\ln w_{Hj} = \ln \frac{\theta b}{1 + \beta_H} + \ln(1 + \beta_H R_j'),$$

where $R_j' = [(\varepsilon - 1)/\varepsilon]v_j/b$. Note that as $\varepsilon \rightarrow \infty$, these reduce to equations (12) and (13). Moreover, regardless of the value of ε , if $\beta_L \approx \beta_H$, then relative wages are constant across firms, and the AKM model of the wage structure remains valid, with the firm effects being monotone functions of value added per worker.

The implied elasticity of wages of skill group S with respect to value added per standardized unit of labor is

$$\xi_{sj} = \frac{\partial \ln w_{sj}}{\partial \ln v_j} = \frac{\beta_s R'_j}{1 + \beta_s R'_j}.$$

Assuming that this elasticity is approximately 0.10 suggests that $\beta_s R'_j \approx 0.10$. Moreover, the elasticity of labor supply of skill group S to the firm is

$$e_{sj} = \frac{\beta_L w_{sj}}{w_{sj} - b_s} = \frac{1 + \beta_s R'_j}{R'_j - 1},$$

so calibrating this elasticity to a value of 4 would suggest that $R'_j = 1.28$, again pointing to a value of $\beta_s \approx 0.08$. Finally, note that the elasticity of employment of skill group S with respect to a change in v_j is

$$e_{sj} \xi_{sj} = \frac{\beta_s R'_j}{R'_j - 1},$$

which has a value of approximately 4 under the preceding assumptions.

When the firm faces a downward-sloping product demand, value added per efficiency unit of labor (v_j) depends on the endogenous choice of output. In the appendix, we show that the elasticities of v_j with respect to an exogenous shift in output demand (indexed by P_j^o) or an exogenous increase in productivity (indexed by T_j) are

$$\begin{aligned} \frac{\partial \ln v_j}{\partial \ln P_j^o} &= \frac{\varepsilon}{\varepsilon + m_j}, \\ \frac{\partial \ln v_j}{\partial \ln T_j} &= \frac{\varepsilon - 1}{\varepsilon + m_j}, \end{aligned}$$

where

$$m_j \equiv \frac{\partial \ln N_j}{\partial \ln v_j} = \frac{R'_j}{R'_j - 1} \left[\beta_L (1 - \theta) \left(\frac{L_j}{N_j} \right) + \beta_H \theta \frac{H_j}{N_j} \right]$$

measures the rate at which overall efficiency units of labor expand when there is an exogenously driven increase in value added. From these expressions, it follows that the elasticities of the wages of skill group S with respect to demand shocks and productivity shocks are

$$\begin{aligned} \frac{\partial \ln w_{sj}}{\partial \ln P_j^o} &= \frac{\varepsilon}{\varepsilon + m_j} \times \xi_{sj}, \\ \frac{\partial \ln w_{sj}}{\partial \ln T_j} &= \frac{\varepsilon - 1}{\varepsilon + m_j} \times \xi_{sj}. \end{aligned}$$

Under the calibrations above, m_j is approximately 4. Assuming that the firm-specific product demand elasticity is between 3 and 10, the elasticity of wages with respect to a shift in the firm’s demand curve will be between 0.04 and 0.07, and the elasticity with respect to a shift in technological efficiency will be between 0.035 and 0.065.

E. Imperfect Substitution between Skill Groups

A limitation of our baseline model is that it assumes perfect substitutability between the two skill groups. We now extend the model by assuming that the firm’s output is a CES aggregate of high- and low-skilled labor:

$$Y_j = T_j N_j = T_j f(L_j, H_j),$$

$$f(L_j, H_j) = [(1 - \theta)L_j^\rho + \theta H_j^\rho]^{1/\rho}, \tag{19}$$

where $\rho \in (-\infty, 1]$ and $\sigma = (1 - \rho)^{-1}$ are the elasticity of substitution between the types of labor. The marginal productivities of the two groups take the form

$$T_{jfL} = T_j(1 - \theta)L_j^{\rho-1}N_j^{1-\rho},$$

$$T_{jfH} = T_j\theta H_j^{\rho-1}N_j^{1-\rho}.$$

Assuming that the firm faces a constant price P_j^o for its output, that $b_L = (1 - \theta)b$, and that $b_H = \theta b$, the first-order conditions (10) and (11) evaluate to

$$w_{Lj} = \frac{b(1 - \theta)}{1 + \beta_L} \left[1 + \beta_L R_j \left(\frac{L_j}{N_j} \right)^{-1/\sigma} \right], \tag{20}$$

$$w_{Hj} = \frac{b\theta}{1 + \beta_H} \left[1 + \beta_H R_j \left(\frac{H_j}{N_j} \right)^{-1/\sigma} \right], \tag{21}$$

where $R_j = T_j P_j^o / b = Y_j P_j^o / b N_j = v_j / b$ is value added per standardized unit of labor relative to the reference wage. These differ from the corresponding equations with a linear technology (eqq. [12], [13]) by the terms $(L_j/N_j)^{-1/\sigma}$ and $(H_j/N_j)^{-1/\sigma}$, which adjust the marginal productivities of L and H workers on the basis of their relative employment shares. These terms disappear when $L_j = H_j$ or when σ is large.²⁶

Note that even when $\beta_L = \beta_H$, log wages in this model are nonseparable in worker and firm heterogeneity. This is because firm heterogeneity in fac-

²⁶ Defining $b_j = H_j/L_j$, note that $L_j/N_j = [1 - \theta(1 - b_j^\rho)]^{-1/\rho}$ and $H_j/N_j = b_j[1 - \theta(1 - b_j^\rho)]^{-1/\rho}$. Thus, when $b_j = 1$, $L_j/N_j = H_j/N_j = 1$.

tor proportions leads to firm-specific wage-skill gaps of the sort usually analyzed at the market level (e.g., Katz and Murphy 1992). If skill types were observable, it would be natural to estimate such a model via nonlinear least squares using data on firm value added. With unobserved skill types, an interactive fixed effects specification would be required that allows the firm effects to depend on the unobserved skill ratio at the firm.

To derive the rent-sharing elasticities in this model, we define

$$\tau_{Lj} = \frac{\beta_L R_j (L_j/N_j)^{-1/\sigma}}{1 + \beta_L R_j (L_j/N_j)^{-1/\sigma}},$$

$$\tau_{Hj} = \frac{\beta_H R_j (H_j/N_j)^{-1/\sigma}}{1 + \beta_H R_j (H_j/N_j)^{-1/\sigma}}.$$

These are the elasticities of wages with respect to v_j , ignoring any adjustment to the relative input of L and H labor. They also represent the proportional wage premiums for L and H workers associated with working at a firm with $R = R_j$ relative to a marginal firm with R close to 1. With this notation, we show in the appendix that the elasticities of wages with respect to value added per labor input can be expressed as

$$\xi_{Lj} \equiv \frac{\partial \ln w_{Lj}}{\partial \ln v_j} = \frac{\tau_{Lj} [1 + (\tau_{Hj} e_{Hj} / \sigma)]}{1 + (1/\sigma) [(1 - \kappa_j) \tau_{Lj} e_{Lj} + \kappa_j \tau_{Hj} e_{Hj}]},$$

$$\xi_{Hj} \equiv \frac{\partial \ln w_{Hj}}{\partial \ln v_j} = \frac{\tau_{Hj} [1 + (\tau_{Lj} e_{Lj} / \sigma)]}{1 + (1/\sigma) [(1 - \kappa_j) \tau_{Lj} e_{Lj} + \kappa_j \tau_{Hj} e_{Hj}]},$$

where (as above) e_{Lj} and e_{Hj} are the elasticities of labor supply of L and H workers to the firm and

$$\kappa_j \equiv \frac{(1 - \theta)L_j^\theta}{(1 - \theta)L_j^\theta + \theta H_j^\theta} = \frac{\partial \ln f}{\partial \ln L_j} = 1 - \frac{\partial \ln f}{\partial \ln H_j}.$$

Notice that

$$\lim_{\sigma \rightarrow \infty} \xi_{sj} = \frac{\beta_s R_j}{1 + \beta_s R_j},$$

which is the expression derived above for our baseline case with a linear technology. With imperfect substitution between groups, the value-added elasticities of the two skill groups, ξ_{Lj} and ξ_{Hj} , will depend on τ_{Lj} and τ_{Hj} and on the labor supply elasticities of the two groups.

The general effect of allowing imperfect substitution between L and H labor is to generate value added elasticities that are smaller than the rent

shares τ_{Lj} and τ_{Hj} but of similar magnitudes. For example, consider a calibration with $\theta = 0.6$, $\sigma = 1.4$, and an average workforce comprised of 60% L workers and 40% H workers. In this setup, a value of $R_j = 1.45$ combined with $\beta_L = 0.09$ and $\beta_H = 0.15$ yields $\tau_L = 0.10$ and $\tau_H = 0.20$, implying that lower-skilled workers earn a 10% wage premium relative to their outside wage while higher-skilled workers earn a 20% premium. The implied firm-specific labor supply elasticities of the two groups are 5 and 2, respectively, while the implied elasticities of their wages with respect to value added per worker are 0.05 and 0.11, respectively.

To summarize, a CES technology will give rise to generalized AKM models in which the wage differentials between different skill groups vary across firms, and the elasticity of wages with respect to value added per worker will also vary across groups and across firms, depending on the firm-specific skill employment shares at the firm. Nevertheless, the basic intuition of our benchmark model remains. In particular, a simple model of monopsonistic wage setting can explain the existence and quantitative magnitude of systematic firm wage premiums that are highly correlated across skill groups and highly correlated with measures of firm-specific productivity.

F. Relationship to Other Models and Open Questions

Although we have worked with a static model of employer differentiation, there are obvious benefits to considering more realistic dynamic models, not least of which is that they provide a rationale for the worker mobility typically used to estimate firm wage effects. Section A2 considers a simple dynamic extension of our framework that yields random mobility between firms and has essentially identical steady-state implications for wages and employment. However, it would be interesting to consider richer models where workers systematically climb a productivity job ladder and can spend some time unemployed. Another interesting extension would be to allow incumbent workers to face switching costs that lead firms to price discriminate against them. This could lead to offer-matching behavior (as in Postel-Vinay and Robin 2002) and to new predictions about recruitment and retention policies.

By assuming that the number of employers is very large, we have adopted a partial equilibrium framework with no strategic interactions between employers. With a finite number of firms, a shock to one firm's productivity will affect the equilibrium employment and wages of competitor firms. Staiger, Spetz, and Phibbs (2010) provide compelling evidence of such responses in the market for nurses. As in the oligopoly literature, analysis of a finite employer model with strong strategic dependence may be complicated by the presence of multiple equilibria, which requires different methods for estimation (e.g., Ciliberto and Tamer 2009) but may also yield interesting policy implications.

Finally, it is worth noting some links between our modeling of workplace differentiation with the literature on compensating differentials for non-

wage amenities (Rosen 1986; Hwang et al. 1998). In our model, nonwage amenities that are valued equally by all workers simply shift the intercept of the labor supply curve to the firm. But a monopsonist firm sets wages on the basis of the elasticity of labor supply to the firm, which is governed entirely by the distribution of taste heterogeneity. For this reason, our model exhibits no compensating differentials of the standard sort. Amenities affect firm effects only through their influence on TFP: a firm with attractive nonwage amenities will grow large, which should depress its revenue productivity and therefore lower its firm wage effect. Empirically distinguishing this effect, which is mediated through product prices, from the standard compensation mechanism is policy relevant, since the monopsony model will tend to imply a different incidence of, for example, employer-provided health benefits on workers than a compensating differentials model.

VI. Conclusions

There is no doubt that a large share of wage inequality is driven by differences in worker skills. But economists have long had evidence (e.g., Lester 1946; Slichter 1950) that employer characteristics exert an independent effect on wages. While the ability of firms to set wages is disciplined by market competition, there are clearly limits to those competitive forces, which also fail to eliminate productivity and output price differences across firms (Hsieh and Klenow 2009).

Modern search theory provides one rationale for the existence of wage-setting power (Mortensen 2005). But even without search frictions, firms will be able to set wages if (as seems likely) workers differ in their valuation of firms' nonwage characteristics. While the mechanisms giving rise to market power under these two approaches are different, both imply that labor is supplied inelastically to firms, providing some scope to set wages. As noted by Manning (2011), such market power 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 setting, firms set wages unilaterally and rents are shared only because of information asymmetries. We believe that this alternative explanation for what the literature has called rent sharing is more plausible than one based on worker bargaining power, particularly in economies that lack strong unions and in settings (like Portugal) where most firms pay wages above those required by sectoral bargaining agreements. Indeed, structural estimates of worker bargaining strength are often quite low (e.g., Cahuc et al. 2006). One interesting implication of a monopsony-based explanation for the link between productivity and wages is that there is no holdup problem in the firm's investment decision, a prediction consistent with the empirical findings of Card et al. (2014).

The empirical literature on firm wage inequality has progressed dramatically with the introduction of huge matched employer-employee data sets.

Yet significant challenges remain. The field continues to rely almost exclusively on observational studies predicated on plausible—but ultimately debatable—identifying assumptions. More research is needed using research designs that can credibly identify the causal link from firm-specific shocks to workers' wages. Another outstanding goal is the development of studies that directly manipulate incentives for workers to leave and join particular firms, as in the innovative experimental design of Dal Bó, Finan, and Rossi (2013). Such designs can be used to rigorously assess the degree of bias in observational firm-switching designs.

While research on labor market inequality typically strives for general explanations of national trends, the way forward in this literature may not involve a theory of everything but rather more attention to the institutional details of particular labor markets. This is the tradition in the industrial organization literature, where studies generally focus on particular industries rather than the economy as a whole. It is plausible that firms have more wage-setting power in some labor markets than others and that the nature of firm wage versus nonwage competition differs as well. A key empirical problem is how to define the labor market of interest both geographically and with respect to skills. In the immigration literature, for example, there is debate over whether labor markets are effectively national or local and over how to classify different age, education, and national origin groups.²⁷ Manning and Petrongolo (2011) develop a spatial job search model where markets can be geographically overlapping. Fitting their model to spatially detailed data on job applications and vacancies, they find that workers are discouraged from searching in areas with strong competition from other job seekers and that shocks to local neighborhoods can yield important ripple effects on labor market activity in nearby areas.²⁸ New case studies of settings where the market structure of labor demand can be carefully documented would be particularly useful.²⁹

²⁷ For example, Borjas (2003) argues for national labor markets categorized into five education classes, with no distinction between immigrants and natives, whereas Card (2009) and Ottaviano and Peri (2012) argue for city-level markets categorized into two education classes but stratified by immigrant origin.

²⁸ Recent research on trade shocks (e.g., Autor, Dorn, and Hansen 2013) and the effects of the Great Recession (e.g., Yagan 2016) shows surprising large and persistent effects of localized shocks, consistent with the idea that the elasticity of supply to local labor markets is relatively inelastic.

²⁹ For example, the empirical literature on monopsony has focused on the market for nurses (e.g., Staiger et al. 2010) and teachers (e.g., Ransom and Sims 2010; Falch 2011) on the basis of a presumption that firms have more wage-setting power in these occupational labor markets. By contrast, Ashenfelter and Hannan (1986) and Black and Strahan (2001) use the product market to define their labor market of interest: they study the effects of banking deregulation on the gender composition and relative wages of bank employees.

Finally, the idea that even highly advanced labor markets, like that of the United States, might be better characterized as imperfectly competitive opens a host of questions about the welfare implications of industrial policies and labor market institutions, such as the minimum wage, unemployment insurance, and employment protection (Katz and Summers 1989; Acemoglu 2001; Coles and Mortensen 2016). Empirical work lags particularly far behind the theory in this domain. Additional evidence on how actual labor market policies affect firm and worker behavior is needed to assess the plausibility of these theoretical policy arguments.

Appendix

A1. Derivations

A1.1. Downward-Sloping Product Demand

Let $\kappa_j = \partial \ln N_j / \partial \ln L_j = (1 - \theta)L_j/N_j$ represent the elasticity of total labor efficiency units with respect to low-skilled labor, and notice that the elasticity of labor inputs with respect to high-skilled labor is $\partial \ln N_j / \partial \ln H_j = \theta H_j/N_j = 1 - \kappa_j$. Suppose that the elasticities of wages of the two groups with respect to value added per worker are ξ_{Lj} and ξ_{Hj} , respectively, and that e_L and e_H are the elasticities of labor supply of the two groups to the firm. Since the firm’s labor input choices are constrained by the labor supply functions of L and H labor, the elasticity of total labor input with respect to a shift in value added per worker is

$$\begin{aligned}
 m_j &\equiv \frac{\partial \ln N_j}{\partial \ln v_j} = \frac{\partial \ln N_j}{\partial \ln L_j} \times \frac{\partial \ln L_j}{\partial \ln w_{Lj}} \times \frac{\partial \ln w_{Lj}}{\partial \ln v_j} + \frac{\partial \ln N_j}{\partial \ln H_j} \times \frac{\partial \ln H_j}{\partial \ln w_{Hj}} \times \frac{\partial \ln w_{Hj}}{\partial \ln v_j} \\
 &= \kappa_j e_{Lj} \xi_{Lj} + (1 - \kappa_j) e_{Hj} \xi_{Hj}.
 \end{aligned}$$

Now, using the fact that value added per worker is $v_j = P_j^\circ T_j^{1-1/\epsilon} N_j^{-1/\epsilon}$, it follows that

$$\begin{aligned}
 \frac{\partial \ln v_j}{\partial \ln P_j^\circ} &= 1 - \frac{1}{\epsilon} m_j \frac{\partial \ln v_j}{\partial \ln P_j^\circ} \\
 \Rightarrow \frac{\partial \ln v_j}{\partial \ln P_j^\circ} &= \frac{\epsilon}{\epsilon + m_j},
 \end{aligned}$$

and similarly

$$\frac{\partial \ln v_j}{\partial \ln T_j} = \frac{\epsilon - 1}{\epsilon + m_j}.$$

A1.2. Constant Elasticity of Substitution (CES) Technology

We now extend the model by assuming that the firm’s production f is in the CES class, so the labor input at firm j is

$$N_j = f(L_j, H_j) = [(1 - \theta)L_j^\rho + \theta H_j^\rho]^{1/\rho}.$$

As noted in the text, the marginal products of the two skill groups are

$$f_L = (1 - \theta)L_j^{\rho-1}f(L_j, H_j)^{1-\rho},$$

$$f_H = \theta H_j^{\rho-1}f(L_j, H_j)^{1-\rho}.$$

Finally, define

$$\kappa_j \equiv \frac{\partial \ln N_j}{\partial \ln L_j} = \frac{(1 - \theta)L_j^\rho}{(1 - \theta)L_j^\rho + \theta H_j^\rho}$$

and note that $\partial \ln N_j / \partial \ln H_j = 1 - \kappa_j$.

The first-order conditions (20) and (21) can be written as

$$\ln \omega_{L_j} = \ln \frac{b(1 - \theta)}{1 + \beta_L} + \ln \left(1 + \beta_L R_j \left(\frac{L_j}{N_j} \right)^{-1/\sigma} \right),$$

$$\ln \omega_{H_j} = \ln \frac{b\theta}{1 + \beta_H} + \ln \left(1 + \beta_H R_j \left(\frac{H_j}{N_j} \right)^{-1/\sigma} \right).$$

Differentiating these equations and simplifying notation, we obtain

$$\begin{bmatrix} 1 + \frac{1}{\sigma} \tau_L (1 - \kappa_j) e_L & -\frac{1}{\sigma} \tau_L (1 - \kappa_j) e_H \\ -\frac{1}{\sigma} \tau_H \kappa_j e_L & 1 + \frac{1}{\sigma} \tau_H \kappa_j e_H \end{bmatrix} \begin{bmatrix} d \ln \omega_{L_j} \\ d \ln \omega_{H_j} \end{bmatrix} = \begin{bmatrix} \tau_L \\ \tau_H \end{bmatrix} dv_j.$$

Some manipulation establishes that

$$\frac{\partial \ln \omega_{L_j}}{\partial \ln v_j} = \frac{\tau_{L_j} [1 + (\tau_{H_j} e_{H_j} / \sigma)]}{1 + (1/\sigma) [(1 - \kappa_j) \tau_{L_j} e_{L_j} + \kappa_j \tau_{H_j} e_{H_j}]},$$

$$\frac{\partial \ln \omega_{H_j}}{\partial \ln v_j} = \frac{\tau_{H_j} [1 + (\tau_{L_j} e_{L_j} / \sigma)]}{1 + (1/\sigma) [(1 - \kappa_j) \tau_{L_j} e_{L_j} + \kappa_j \tau_{H_j} e_{H_j}]}.$$

A2. Two-Period Model of Supply

Here we consider a two-period extension of our static framework. A worker i of type S faces indirect utility over firms $j \in \{1, \dots, J\}$ of

$$u_{is_j} = \beta_S \ln(\omega_{Sj} - b_S) + a_{Sj} + \epsilon_{is_j},$$

where ϵ_{is_j} is drawn from a type I extreme value distribution. Hence, the period 1 choice probabilities are

$$p_{sj}^1 = \frac{\exp(\beta_s \ln(w_{sj}^1 - b_s) + a_{sj}^1)}{\sum_{k=1}^J \exp(\beta_s \ln(w_{sk}^1 - b_s) + a_{sk}^1)},$$

$$\approx \lambda_s^1 \exp(\beta_s \ln(w_{sj}^1 - b_s) + a_{sj}^1).$$

In the second period, a fraction $\tilde{\pi}$ of the workers get a new draw ϵ'_i of idiosyncratic extreme value preferences. Because each firm’s market share is very low, workers will choose only employers for which they have a very strong idiosyncratic taste. Hence, the chances of preferring to stay at the same firm with a new taste draw are essentially zero. With this in mind, we write second-period market shares as

$$p_{sj}^2 = \tilde{\pi} \frac{\exp(\beta_s \ln(w_{sj}^2 - b_s) + a_{sj}^2)}{\sum_{l=1}^J \exp(\beta_s \ln(w_{sl}^2 - b_s) + a_{sl}^2)} + (1 - \tilde{\pi})p_{sj}^1$$

$$\approx \tilde{\pi} \lambda_s^2 \exp(\beta_s \ln(w_{sj}^2 - b_s) + a_{sj}^2) + (1 - \tilde{\pi})p_{sj}^1.$$

Clearly, as $\tilde{\pi} \rightarrow 1$, the labor supply function becomes static. Otherwise, we have a partial adjustment process that yields heterogeneity in the labor supply elasticity, depending on how far p_{sj}^1 is from $\lambda_s^2 \exp(\beta_s \ln(w_{sj}^2 - b_s) + a_{sj}^2)$. In a steady state these two objects will be the same, and the elasticity of supply to each firm simplifies to $\tilde{\pi}$ times the usual static elasticity e_{sj} .

Therefore, we can think about the steady state of a dynamic model with taste shocks as one where firms face a supply curve with elasticity $e_{sj}\tilde{\pi}$ and set wages accordingly. As before, firms cannot observe workers’ preferences. Hence, employee threats to leave in response to taste shocks will not be viewed as credible by the firm, despite the firm’s knowledge that a fraction $\tilde{\pi}$ of workers did in fact draw new tastes. Because the firm cannot budge in its wage policy, each period will yield a fraction $\tilde{\pi}$ of workers switching between firms.

Table A1
Summary of Estimated Rent-Sharing Elasticities

Study	Design Features	Measure of Profitability	Elasticity
A. Industry-level profit measures:			
Christofides and Oswald 1992	Canadian union contracts; 120 narrowly defined manufacturing industries	Industry profits/worker (wage changes)	.07
Blanchflower et al. 1996	US individual wage data (CPS), grouped to industry \times year cells; manufacturing only	Industry profits/worker (within-industry changes)	.01–.06

Table A1 (Continued)

Study	Design Features	Measure of Profitability	Elasticity
A. Industry-level profit measures:			
Estevao and Tevlin 2003	US manufacturing industry data; adjusted for labor quality; instrument for value added = demand shocks in downstream sectors	Value added per worker (first differences)	.29
		Profit per worker (first differences)	.14
B. Firm-level profit measures, average firm-level wages:			
Abowd and Lemieux 1993	Canadian union contracts merged to corporate accounts; instruments for revenues = industry selling prices, import and export prices	Quasi rent/worker (wage change model)	.22
Van Reenen 1996	Large British manufacturing firms merged with corporate accounts; instruments for rents = innovations, imports, research and development, industry concentration	Quasi rent/worker (wage change model)	.29
Hildreth and Oswald 1997	British firms (EXSTAT); firm-specific profits (from financial statements); instruments = lagged values of wages and profits	Profit per worker	.02
Hildreth 1998	British manufacturing establishments; establishment-specific value added; instruments for rents = innovation measure	Quasi rent/worker	.03
Barth et al. 2016	US establishments in LBD; establishment-specific revenues; instrument for revenues/worker = revenues/worker in same industry, other regions	Sales/worker (within-establishment changes)	.32 (OLS) .16 (IV)

Table A1 (Continued)

Study	Design Features	Measure of Profitability	Elasticity
C. Individual wages and firm-level profit measures:			
Margolis and Salvanes 2001	Worker and firm data for France and Norway; full-time male workers in manufacturing; profit from financial filings; instruments = sales/worker and subsidies/worker	Profit per worker	.03 (France) .01 (Norway)
Arai 2003	Swedish worker panel matched to employer (10-year stayers design); profits from financial statements	Change in 5-year average profit per worker	.01-.02
Guiso et al. 2005	Italian worker panel matched to larger firms; value added from financial statements; model-based decomposition of value-added shocks	Permanent shock to log value added per worker	.07
		Transitory shock to log value added per worker	.00
Fakhfakh and FitzRoy 2004	Larger French manufacturing establishments; value added from establishment survey	Mean log value added/worker over past 3 years	.12
Du Caju, Rycx, and Tojerow 2011	Belgian establishment panel; value added and labor cost from financial statements	Value added minus labor costs per worker	.03-.04
Martins 2009	Larger Portuguese manufacturing firms; revenue and capital costs from financial statements; instruments = export share of sales \times exchange rate changes	Revenue-capital costs/worker (differenced)	.03-.05
Gürtzgen 2009	German establishment/worker panel (LIAB) value added from establishment survey; instruments for change in quasi rent = lags of value added and wages	Quasi rent/worker (no adjustment for capital)	.03-.04
		Change in quasi rent/worker (stayers design)	.01-.06

Table A1 (Continued)

Study	Design Features	Measure of Profitability	Elasticity
Cardoso and Portela 2009	Portuguese worker panel; sales from firm reports; model-based decomposition of sales shocks	Permanent shock to log sales	.09
		Transitory shock to log sales	.00
Arai and Heyman 2009	Swedish worker/firm panel; profits from financial statements; stayers design; instrument = change in foreign sales	Change in profit per worker	.07
Card et al. 2014	Italian worker panel matched to firms; value added and capital from financial statements; instrument for value added = sales/worker at firms in other regions	Value added per worker (within job match)	.06–.08
Carlsson et al. 2014	Swedish worker panel matched to firms; mining and manufacturing only; firm-specific output and selling price indexes; instruments for productivity = indexes of firm-specific and sectoral TFPQ	Firm-specific output/worker (within job match)	.05
		Sectoral average output/worker (within job match)	.15
Card et al. 2016	Portuguese worker panel matched to firms; value added and capital from financial statements; wage measure = estimated firm effect from AKM model	Mean value added per worker	.16 (males) .14 (females)
		Mean value added per worker (changes for stayers)	.05 (males) .04 (females)
Bagger et al. 2014	Danish worker panel matched to firms; output from firm survey; nonparametric regressions within sector of wages on labor productivity	Output per worker	.09 (manufacturing) .13 (trade) .05 (transp./comm.) .07 (finance/real estate)

NOTE.—Estimates were extracted by authors from studies listed. AKM = Abowd, Kramarz, and Margolis (1999); CPS = Current Population Survey; EXSTAT = Exstat database; IV = instrumental variables; LBD = Longitudinal Business Database; LIAB = Linked Employer-Employee Data of the Institute for Employment Research; OLS = ordinary least squares; TFPQ = physical total factor productivity.

References

- Abowd, John M., Francis Kramarz, Paul A. Lengeremann, and Sébastien Pérez-Duarte. 2004. Are good workers employed by good firms? A test of a simple assortative matching model for France and the United States. Unpublished manuscript, Centre de Recherche en Economie et en Statistique.
- Abowd, John M., Franic Kramarz, and David N. Margolis. 1999. High wage workers and high wage firms. *Econometrica* 67, no. 2:251–333.
- Abowd, John M., and Thomas Lemieux. 1993. The effects of product market competition on collective bargaining agreements: The case of foreign competition in Canada. *Quarterly Journal of Economics* 108, no. 4:983–1014.
- Abowd, John M., Paul A. Lengeremann, and Kevin L. McKinney. 2003. The measurement of human capital in the U.S. economy. Technical Paper no. TP-2002-09, Center for Economic Studies, US Census Bureau, Washington, DC.
- Acemoglu, Daron. 2001. Good jobs versus bad jobs. *Journal of Labor Economics* 19, no. 1:1–21.
- Acemoglu, Daron, and David H. Autor. 2011. Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of labor economics*, vol. 4, ed. David Card and Orley Ashenfelter, 1043–71. Amsterdam: Elsevier.
- Akerlof, George A., and Janet L. Yellen. 1986. *Efficiency wage models of the labor market*. Cambridge: Cambridge University Press.
- Albrecht, James W., and Bo Axell. 1984. An equilibrium model of search unemployment. *Journal of Political Economy* 92, no. 5:824–40.
- Andrews, Martyn J., Leonard Gill, Thorsten Schank, and Richard Upward. 2008. High wage workers and low wage firms: Negative assortative matching or limited mobility bias? *Journal of the Royal Statistical Society: Series A* 171, no. 3:673–97.
- . 2012. High wage workers match with high wage firms: Clear evidence of the effects of limited mobility bias. *Economics Letters* 117, no. 3: 824–27.
- Arai, Mahmood. 2003. Wages, profits, and capital intensity: Evidence from matched worker-firm data. *Journal of Labor Economics* 21, no. 3:593–618.
- Arai, Mahmood, and Fredrik Heyman. 2009. Microdata evidence on rent-sharing. *Applied Economics* 41, no. 23:2965–76.
- Ashenfelter, Orley, and Timothy Hannan. 1986. Sex discrimination and product market competition: The case of the banking industry. *Quarterly Journal of Economics* 101, no. 1:149–73.
- Autor, David H., David Dorn, and Gordon H. Hanson. 2013. The China syndrome: local labor market effects of import competition in the United States. *American Economic Review* 103, no. 6:2121–68.

- Bagger, Jesper, Bent J. Christensen, and Dale T. Mortensen. 2014. Wage and labor productivity dispersion: The roles of total factor productivity, labor quality, capital intensity, and rent sharing. Mimeo, Royal Holloway.
- Barth, Erling, Alex Bryson, James C. Davis, and Richard B. Freeman. 2016. It's where you work: Increases in the dispersion of earnings across establishments and individuals in the United States. *Journal of Labor Economics* 34, no. S2:S67–S97.
- Beaudry, Paul, David A. Green, and Benjamin M. Sand. 2012. Does industrial composition matter for wages? A test of search and bargaining theory. *Econometrica* 80, no. 3:1063–104.
- Bernard, Andrew B., and J. Bradford Jensen. 1995. Exporters, jobs, and wages in US manufacturing: 1976–1987. *Brookings Papers on Economic Activity: Microeconomics* 1995:67–119.
- Berry, Steven T. 1994. Estimating discrete-choice models of product differentiation. *RAND Journal of Economics* 25, no. 2:242–62.
- Berry, Steven T., James Levinsohn, and Ariel Pakes. 1995. Automobile prices in market equilibrium. *Econometrica* 63, no. 4:841–90.
- Berry, Steven T., and Ariel Pakes. 2007. The pure characteristics demand model. *International Economic Review* 48, no. 4:1193–1225.
- Black, Sandra E., and Philip E. Strahan. 2001. The division of spoils: rent-sharing and discrimination in a regulated industry. *American Economic Review* 91, no. 4:814–31.
- Blanchflower, David G., Andrew J. Oswald, and Peter Sanfey. 1996. Wages, profits and rent-sharing. *Quarterly Journal of Economics* 111, no. 1:227–51.
- Boal, William M., and Michael R. Ransom. 1997. Monopsony in the labor market. *Journal of Economic Literature* 35, no. 1:86–112.
- Bonhomme, Stéphane, and Gregory Jolivet. 2009. The pervasive absence of compensating differentials. *Journal of Applied Econometrics* 24, no. 5: 763–95.
- Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa. 2015. A distributional framework for matched employer employee data. Unpublished manuscript, University of Chicago.
- Borjas, George B. 2003. The labor demand curve is downward sloping: Re-examining the impact of immigration on the labor market. *Quarterly Journal of Economics* 118, no. 4:1335–74.
- Burdett, Kenneth, and Dale T. Mortensen. 1998. Wage differentials, employer size, and unemployment. *International Economic Review* 39, no. 2:257–73.
- Cahuc, Pierre, Fabien Postel-Vinay, and Jean-Marc Robin. 2006. Wage bargaining with on-the-job search: Theory and evidence. *Econometrica* 74, no. 2:323–64.
- Card, David. 2009. Immigration and inequality. *American Economic Review* 90, no. 2:1–21.

- Card, David, and Ana Rute Cardoso. 2012. Can compulsory military service raise civilian wages? Evidence from the peacetime draft in Portugal. *American Economic Journal: Applied Economics* 4, no. 4:57–93.
- Card, David, Ana Rute Cardoso, and Patrick M. Kline. 2016. Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women. *Quarterly Journal of Economics* 131, no. 2: 633–86.
- Card, David, Francesco Devicienti, and Agata Maida. 2014. Rent-sharing, holdup, and wages: Evidence from matched panel data. *Review of Economic Studies* 81, no. 1:84–111.
- Card, David, Jörg Heining, and Patrick M. Kline. 2012. Workplace heterogeneity and the rise of West German wage inequality. NBER Working Paper no. 18522, National Bureau of Economic Research, Cambridge, MA.
- . 2013. Workplace heterogeneity and the rise of West German wage inequality. *Quarterly Journal of Economics* 128, no. 3:967–1015.
- Cardoso, Ana Rute. 1997. Workers or employers: Who is shaping wage inequality? *Oxford Bulletin of Economics and Statistics* 59, no. 4:523–47.
- . 1999. Firms' wage policies and the rise in labor market inequality: The case of Portugal. *Industrial and Labor Relations Review* 53, no. 1: 87–102.
- Cardoso, Ana Rute, and Miguel Portela. 2009. Micro foundations for wage flexibility: Wage insurance at the firm level. *Scandinavian Journal of Economics* 111, no. 1:29–50.
- Carlsson, Mikael, Julián Messina, and Oscar N. Skans. 2014. Wage adjustment and productivity shocks. *Economic Journal* 126, no. 595:1739–73, doi: 10.1111/eoj.12214.
- Charles, Kerwin K., and Jonathan Guryan. 2008. Prejudice and wages: An empirical assessment of Becker's *The economics of discrimination*. *Journal of Political Economy* 116, no. 5:773–809.
- . 2011. Studying discrimination: Fundamental challenges and recent progress. *Annual Review of Economics* 3, no. 1:479–511.
- Christofides, Louis N., and Andrew J. Oswald. 1992. Real wage determination and rent-sharing in collective bargaining agreements. *Quarterly Journal of Economics* 107, no. 3:985–1002.
- Ciliberto, Federico, and Elie Tamer. 2009. Market structure and multiple equilibria in airline markets. *Econometrica* 77, no. 6:1791–1828.
- Coles, Melvyn G., and Dale T. Mortensen. 2016. Equilibrium labor turnover, firm growth and unemployment. *Econometrica* 84, no. 1:347–63.
- Dal Bó, Ernesto, Frederico Finan, and Martin A. Rossi. 2013. Strengthening state capabilities: The role of financial incentives in the call to public service. *Quarterly Journal of Economics* 128, no. 3:1169–1218.
- Davis, Steven J., and John Haltiwanger. 1991. Wage dispersion between and within US manufacturing plants, 1963–86. *Brookings Papers on Economic Activity: Microeconomics* 1991:115–200.

- de Menil, Georges. 1971. *Bargaining: Monopoly power versus union power*. Cambridge, MA: MIT Press.
- Du Caju, Philip, François Rycx, and Ilan Tojerow. 2011. Inter-industry wage differentials: How much does rent sharing matter? *Manchester School* 79, no. 4:691–717.
- Dunne, Timothy, Lucia Foster, John Haltiwanger, and Kenneth R. Troske. 2004. Wage and productivity dispersion in United States manufacturing: The role of computer investment. *Journal of Labor Economics* 22, no. 2: 397–429.
- Eeckhout, Jan, and Philipp Kircher, P. 2011. Identifying sorting—in theory. *Review of Economic Studies* 78, no. 3:872–906.
- Estevão, Marcello, and Stacey Tevlin. 2003. Do firms share their success with workers? The response of wages to product market conditions. *Economica* 70, no. 280:597–617.
- Faggio, Giulia, Kjell Salvanes, and John Van Reenen. 2010. The evolution of inequality in productivity and wages: Panel data evidence. *Industrial and Corporate Change* 19, no. 6:1919–51.
- Fakhfakh, Fathi, and Felix FitzRoy. 2004. Basic wages and firm characteristics: Rent sharing in French manufacturing. *Labour* 18, no. 4:615–31.
- Falch, Torberg. 2011. Teacher mobility responses to wage changes: Evidence from a quasi-natural experiment. *American Economic Review* 101, no. 3:460–65.
- Foster, Lucia, John Haltiwanger, and Chad Syverson. 2008. Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? *American Economic Review* 98, no. 1:394–425.
- Gaure, Simen. 2014. Correlation bias correction in two-way fixed-effects linear regression. *Stat* 3, no. 1:379–90.
- Gibbons, Robert, and Lawrence F. Katz. 1992. Does unmeasured ability explain inter-industry wage differentials? *Review of Economic Studies* 59, no. 3:515–35.
- Gibbons, Robert, Lawrence F. Katz, Thomas Lemieux, and Daniel Parent. 2005. Comparative advantage, learning, and sectoral wage determination. *Journal of Labor Economics* 23, no. 4:681–724.
- Goldberg, Pinelopi K., and Rebecca Hellerstein. 2013. A structural approach to identifying the sources of local currency price stability. *Review of Economic Studies*, 80, no. 1:175–210.
- Goldin, Claudia D., and Lawrence F. Katz. 2009. *The race between education and technology*. Cambridge, MA: Harvard University Press.
- Goldschmidt, Deborah, and Johannes F. Schmieder. 2015. The rise of domestic outsourcing and the evolution of the German wage structure. NBER Working Paper no. 21366, National Bureau of Economic Research, Cambridge, MA.

- Gorodnichenko, Yuriy, and Oleksandr Talavera. 2017. Price setting in online markets: Basic facts, international comparisons, and cross-border integration. *American Economic Review* 107, no. 1:249–82.
- Groschen, Erica L. 1991. Sources of intra-industry wage dispersion: How much do employers matter? *Quarterly Journal of Economics* 106, no. 3: 869–84.
- Grout, Paul A. 1984. Investment and wages in the absence of binding contracts: A Nash bargaining approach. *Econometrica* 52, no. 2:449–60.
- Guiso, Luigi, Luigi Pistaferri, and Fabiano Schivardi. 2005. Insurance within the firm. *Journal of Political Economy* 113, no. 5:1054–87.
- Gürtzgen, Nicole. 2009. Rent-sharing and collective bargaining coverage: Evidence from linked employer-employee data. *Scandinavian Journal of Economics* 111, no. 2:323–49.
- Hildreth, Andrew K. G. 1998. Rent-sharing and wages: Product demand or technology driven premia? *Economics of Innovation and New Technology* 5, no. 2–4:199–226.
- Hildreth, Andrew K. G., and Andrew J. Oswald. 1997. Rent-sharing and wages: evidence from company and establishment panels. *Journal of Labor Economics* 15, no. 2:318–37.
- Hornstein, Andreas, Per Krusell, and Giovanni Violante. 2011. Frictional wage dispersion in search models: A quantitative assessment. *American Economic Review* 101, no. 7:2873–98.
- Hsieh, Chang-Tai, and Peter J. Klenow. 2009. Misallocation and manufacturing TFP in China and India. *Quarterly Journal of Economics* 124, no. 4: 1403–48.
- Hwang, Hae-shin, Dale T. Mortensen, and W. Reed. 1998. Hedonic wages and labor market search. *Journal of Labor Economics* 16, no. 4: 815–47.
- Katz, Lawrence F. 1986. Efficiency wage theories: A partial evaluation. *NBER Macroeconomics Annual* 1986:235–90.
- Katz, Lawrence F., and David H. Autor. 1999. Changes in the wage structure and earnings inequality. In *Handbook of labor economics*, vol. 3, ed. Orley Ashenfelter and David Card, 1463–1555. Amsterdam: Elsevier.
- Katz, Lawrence F., and Kevin M. Murphy. 1992. Changes in relative wages, 1963–1987: Supply and demand factors. *Quarterly Journal of Economics*, 107, no. 1:35–78.
- Katz, Lawrence F., and Lawrence H. Summers. 1989. Industry rents: Evidence and implications. *Brookings Papers on Economic Activity: Microeconomics* 1989:209–90.
- Krueger, Alan B., and Lawrence H. Summers. 1988. Efficiency wages and the inter-industry wage structure. *Econometrica* 56, no. 2:259–93.
- Lange, Fabien. 2007. The speed of employer learning. *Journal of Labor Economics* 25, no. 1:1–35.

- Lavetti, Kurt, and Ian M. Schmutte. 2016. Estimating compensating wage differentials with endogenous job mobility. Unpublished manuscript, Ohio State University.
- Lester, Richard A. 1946. Wage diversity and its theoretical implications. *Review of Economics and Statistics* 28, no. 3:152–59.
- . 1952. A range theory of wage differentials. *Industrial and Labor Relations Review* 5, no. 4:483–500.
- Lochner, Benjamin, and Bastian Schulz. 2016. Labor market sorting in Germany. CESifo Working Paper no. 6066, Munich.
- Macis, Mario, and Fabiano Schivardi. 2016. Exports and wages: Rent sharing, workforce composition, or returns to skills? *Journal of Labor Economics* 34, no. 4:945–78.
- Manning, Alan. 2011. Imperfect competition in the labor market. In *Handbook of labor economics*, vol. 4, ed. Orley Ashenfelter and David Card, 973–1041. Amsterdam: Elsevier.
- Manning, Alan, and Barbara Petrongolo. 2011. How local are labor markets? Evidence from a spatial job search model. CEPR Discussion Paper no. DP8686, Centre for Economic Policy Research, Washington, DC.
- Maré, David C., and Dean R. Hyslop. 2006. Worker-firm heterogeneity and matching: An analysis using worker and firm fixed effects estimated from LEED. Statistics New Zealand LEED Research Paper.
- Margolis, David N., and Kjell G. Salvanes. 2001. Do firms really share rents with their workers? IZA Discussion Paper no. 330, IZA Institute of Labor Economics, Bonn.
- Martins, Pedro S. 2009. Rent sharing before and after the wage bill. *Applied Economics* 41, no. 17:2133–51.
- Mas, Alexandre, and Amanda Pallais. 2016. Valuing alternative work arrangements. NBER Working Paper no. 22708, National Bureau of Economic Research, Cambridge, MA.
- McFadden, Daniel L. 1973. Conditional logit analysis of qualitative choice behavior. In *Frontiers in econometrics*, ed. P. Zarembka, 105–42. Academic Press: New York.
- Mincer, Jacob. 1974. Schooling, experience, and earnings. New York: National Bureau of Economic Research.
- Mortensen, Dale T. 2005. *Wage dispersion: Why are similar workers paid differently?* Cambridge, MA: MIT Press.
- Murphy, Kevin M., and Robert H. Topel. 1990. Efficiency wages reconsidered: Theory and evidence. In *Advances in the theory and measurement of unemployment*, ed. Yoram Weiss and Gideon Fishelson, 204–40. Basingstoke: Palgrave Macmillan.
- Oaxaca, Ronald L. 1973. Male-female wage differentials in urban labor markets. *International Economic Review* 14, no. 3:693–709.
- Ottaviano, Gianmarco I. P., and Giovanni Peri. 2012. Rethinking the effect of immigration on wages. *Journal of the European Economic Association* 10, no. 1:152–97.

- Pakes, Ariel. 2016. Empirical tools and competition analysis: Past progress and current problems. NBER Working Paper no. 22086, National Bureau of Economic Research, Cambridge, MA.
- Postel-Vinay, Fabien, and Jean-Marc Robin. 2002. Equilibrium wage dispersion with worker and employer heterogeneity. *Econometrica* 70, no. 6: 2295–350.
- Piyapromdee, Suphanit. 2013. Residual wage dispersion with efficiency wages. Unpublished manuscript, University College London.
- Ransom, Michael R., and David P. Sims. 2010. Estimating the firm's labor supply curve in a "new monopsony" framework: Schoolteachers in Missouri. *Journal of Labor Economics* 28, no. 2:331–55.
- Reynolds, Lloyd G. 1946. Wage differences in local labor markets. *American Economic Review* 36, no. 3:366–75.
- Robinson, Joan. 1933. *The economics of imperfect competition*. London: Macmillan.
- Rosen, Sherwin. 1986. The theory of equalizing differences. In *Handbook of labor economics*, vol. 1, ed. Orley Ashenfelter and Richard Layard, 641–92. Amsterdam: Elsevier.
- Skans, Oskar N., Per-Anders Edin, and Bertil Holmlund. 2009. Wage dispersion between and within plants: Sweden 1985–2000. In *The structure of wages: An international comparison*, ed. Edward P. Lazear and Kathryn L. Shaw, 217–60. Chicago: University of Chicago Press.
- Slichter, Sumner H. 1950. Notes on the structure of wages. *Review of Economics and Statistics* 32, no. 1:80–91.
- Song, Jae, David J. Price, Fatih Guvenen, Nicholas Bloom, and Till von Wachter. 2015. Firming up inequality. NBER Working Paper no. 21199, National Bureau of Economic Research, Cambridge, MA.
- Sorkin, Isaac. 2015. Ranking firms using revealed preference. Unpublished manuscript, Stanford University.
- Staiger, Douglas O., Joanne Spetz, and Ciaran S. Phibbs. 2010. Is there monopsony in the labor market? Evidence from a natural experiment. *Journal of Labor Economics* 28, no. 2:211–36.
- Svejnjar, Jan. 1986. Bargaining power, fear of disagreement, and wage settlements: Theory and evidence from US industry. *Econometrica* 54, no. 5: 1055–78.
- Syverson, Chad. 2004. Market structure and productivity: A concrete example. *Journal of Political Economy* 112, no. 6:1181–222.
- . 2011. What determines productivity? *Journal of Economic Literature* 113, no. 3:326–65.
- Tabler, Christopher, and Rune Vejlin. 2016. Estimation of a Roy/search/compensating differential model of the labor market. NBER Working Paper no. 22439, National Bureau of Economic Research, Cambridge, MA.
- Tirole, Jean. 1988. *The theory of industrial organization*. Cambridge, MA: MIT Press.

- Topel, Robert H., and Michael P. Ward. 1992. Job mobility and the careers of young men. *Quarterly Journal of Economics* 107, no. 2:439–79.
- Van Reenen, John. 1996. The creation and capture of rents: Wages and innovation in a panel of UK companies. *Quarterly Journal of Economics* 111, no. 1:195–226.
- Yagan, Danny. 2016. Is the Great Recession really over? Evidence of enduring employment impacts. Unpublished manuscript, University of California Berkeley.