

Firms and Labor Market Inequality: Evidence and Some Theory

David Card, Ana Rute Cardoso, Joerg Heining, and Patrick Kline*

March 2016

Abstract

We review the literature on firm-level drivers of labor market inequality. There is strong evidence from a variety of fields that standard measures of productivity – like output per worker or total factor productivity – vary substantially across firms, even within narrowly-defined industries. Several recent studies note that rising trends in the dispersion of productivity across firms mirror the trends in the wage inequality across workers. Two distinct literatures have searched for a more direct link between these two phenomena. The first examines how wages are affected by differences in employer productivity. Studies that focus on firm-specific productivity shocks and control for the non-random sorting of workers to more and less productive firms typically find that a 10% increase in value-added per worker leads to somewhere between a 0.5% and 1.5% increase in wages. A second literature focuses on firm-specific wage premiums, using the wage outcomes of job changers. This literature also concludes that firm pay setting is important for wage inequality, with many studies finding that firm wage effects contribute approximately 20% of the overall variance of wages. To interpret these findings, we develop a model where workplace environments are viewed as imperfect substitutes by workers, and firms set wages with some degree of market power. We show that simple versions of this model can readily match the stylized empirical findings in the literature regarding rent-sharing elasticities and the structure of firm-specific pay premiums.

*We are extremely grateful to Raffaele Saggio for assistance in preparing this paper, and to David Green for helpful suggestions on an earlier draft.

Does where you work determine how much 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 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 focus 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). Though 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 possibility that wages reflect anything other than skill remains highly controversial.

The growing availability of matched employer-employee datasets has created new opportunities to disentangle the effects of worker and firm heterogeneity on wage inequality. Nevertheless, many of the fundamental issues that economists have long debated about differences in the characteristics of the workers at different firms, and the nature of the jobs at different workplaces, carry over to these new datasets.² This review summarizes what has been learned so far from these new datasets 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 show 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. 2014). 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 differences in firm productivity on the wages of workers. The resulting estimates are typically expressed as “rent-sharing” elasticities. The findings in this literature 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 0.05-0.15, though 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 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 the two-way fixed effects estimator of Abowd, Kramarz, and Margolis (1999) (hereafter AKM), 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. Specifically, we argue that the firm-specific wage premiums estimated in an AKM model

¹This 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).

²Many of the issues about the interpretation of firm-specific wage setting closely parallel issues that were raised in the earlier literature on industry-specific wage premiums – see e.g., Slichter (1950), Katz (1986), Krueger and Summers (1988), Gibbons and Katz (1992), Katz and Summers (1989), and Murphy and Topel (1990).

incorporate any rent-sharing effect, while adjusting for observed or unobserved skill differences between workers at different firms (which are absorbed by the estimated worker effects in these models). Using data from Portugal we show that more productive firms pay higher average wage premiums relative to the outside labor market, 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.

We then go on to investigate the extent of *differential* rent sharing between more and less educated workers in the Portuguese labor market. We confirm that more productive firms have a larger share of highly-educated workers. Nevertheless, the wage premiums offered by more productive firms to more- and less-educated workers are very similar, and the *relative wage* of highly educated workers is nearly constant across firms, consistent with the additive specification underlying the AKM model.

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 labor market states. 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.³ This workplace differentiation could reflect heterogeneity in firm location, job characteristics (e.g., corporate 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. Our model can be viewed as an adaptation of the standard random preferences model of consumer demand (e.g., Berry, 1994; Berry, Levinsohn, and Pakes, 1995), with firms setting wages rather than prices.

We presume, as in Robinson’s analysis and much of the Industrial Organization literature, that the firm cannot price discriminate based upon 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. 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.

Conditions are provided under which log wages are additively separable into components due to worker and firm heterogeneity, as in the pioneering econometric model of AKM. Specifically, we show that the firm-specific wage premium will be constant across skill groups if different groups are perfect substitutes in production, or if different skill groups have similar elasticities of supply to the firm. 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.

We conclude with some thoughts on unresolved empirical and theoretical issues in the literature. Perhaps the most important empirical concern is the lack of quasi-experimental sources of variation in firm-specific productivity or firm switching. While a few older studies attempt to leverage world prices (Abowd and Lemieux, 1993) or product market innovations (Van Reenen, 1996) to identify rent sharing elasticities, most

³In their review of monopsony models, Boal and Ransom (1997) refer to this as the case of “classic differentiation”.

recent studies, while able to control for worker heterogeneity, have not compellingly isolated exogenous changes in productivity. On the theoretical side, an important issue is how far the insights from a simple static wage setting model carry over to frictional labor market settings.

1 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 two within 4-digit industries. Hsieh and Klenow (2009) find even larger productivity gaps in India and China, with 90-10 *TFP* ratios on the order of five. 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 strongly correlated with firm survival (Foster, Haltiwanger, and Syverson, 2008).

It is natural to wonder if 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 (Slichter, 1950; Davis and Haltiwanger, 1991; Groshen, 1991; Bernard and Jensen, 1995; Cardoso, 1997; Cardoso, 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, Salvanes, and Van Reenen, 2010; Barth et al. 2014). Figure 1 plots results from Barth et al. (2014), showing remarkably similar trends in the dispersion of wages and productivity across business establishments in the United States. Taken at face value, the parallel trends are consistent with a roughly unit elasticity of establishment wages with respect to productivity (see Barth et al, 2014, p. 5). 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 standard 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. Appendix Table 1 describes 21 recent studies in this literature. The basic idea in these papers is to relate wages to some measure of employer profitability or rents. Since different studies use different measures of rents, however, it is important to clarify how different choices affect the estimated rent sharing elasticity that is reported in a given study. It is also important to clarify the role of heterogeneity in workers' skills, which can confound estimation.

Measuring rents

For simplicity, we will work with a model with two types of labor, and ignore capital. Define the profits earned by firm j as:

$$\pi_j = VA_j - w_{Lj}L_j - w_{Hj}H_j,$$

where VA_j is value added, L_j and H_j represent employment of “low skill” and “high skill” labor at firm j , and w_{L_j} and w_{H_j} denote the wages paid to the two types of labor. Assume that value added is produced by a linear technology:

$$VA_j \equiv R_j - M_j = P_j T_j ((1 - \theta) L_j + \theta H_j)$$

where R_j represents sales, M_j represents the cost of materials and other intermediate inputs (e.g., energy), P_j is a potentially *firm-specific* selling price index, T_j is an index of technical efficiency, and θ is an index of the relative efficiency of type H workers. Here $P_j T_j$ represents total factor productivity (TFP_j) which, in the terminology of Foster, Haltiwanger and Syverson (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.

Letting $N_j = L_j + H_j$ represent the total number of workers at the firm, value added per worker is $\frac{VA_j}{N_j} = TFP_j q_j$ where $q_j = \frac{(1-\theta)L_j + \theta H_j}{N_j}$ is the average quality of the firm’s workforce. The logarithm of value added per worker is:

$$\ln \left(\frac{VA_j}{N_j} \right) = \ln TFP_j + \ln q_j.$$

Holding constant labor quality, value added per worker is therefore a valid index of TFP . When differences in labor quality are ignored (or imperfectly measured), however, there are two problems with the use of value added per worker as an index of productivity. The first is that average wages at a firm will (in general) depend on average worker quality. A firm with higher quality workers will have higher value added per worker *and* higher average wages, leading to an upward bias in rent sharing models based on firm-wide average wages. The second is that value added per worker is more variable than TFP . This can lead to attenuation bias in specifications that relate wages for a specific subgroup of workers to value added per worker at the firm.

Instead of using value added per worker, some studies use sales per worker as a measure of productivity. Assuming that intermediate inputs vary proportionally with revenues (i.e., $M_j = m_j R_j$), sales per worker can be decomposed as:

$$\ln \left(\frac{R_j}{N_j} \right) = \ln TFP_j + \ln q_j - \ln(1 - m_j),$$

which varies with TFP , labor quality, and the fraction of intermediate inputs in final sales. Sales per worker has the same potential problems as value added per worker, plus the extra complication introduced by variation across firms in the fraction of intermediate inputs and services that are purchased rather than produced in-house.

Many rent sharing studies adopt the bargaining framework laid out by de Menil (1971), in which workers and the firm split a so-called “quasi-rent”:⁴

$$Q_j \equiv VA_j - w_{L_j}^a L_j - w_{H_j}^a H_j,$$

where $(w_{L_j}^a, w_{H_j}^a)$ are the alternative wages available to workers in the event of a breakdown in negotiations. Quasi-rent per worker is $\frac{Q_j}{N_j} = \frac{VA_j}{N_j} - w_{L_j}^a (1 - s_j) - w_{H_j}^a s_j$ where $s_j = \frac{H_j}{N_j}$ gives the fraction of high-skilled

⁴Most studies in the recent literature ignore the determination of employment and also ignore capital. Svejnar (1986) presents an analysis that generalizes de Menil (1971) to allow for the optimal determination of employment. When the firm also has to select a capital stock prior to the determination of wages there is also a potential hold-up problem in the choice of capital (Grout, 1984). Card, Devicienti, and Maida (2014) argue that hold up does not appear to be a major issue for Italian firms.

workers at the firm. The elasticity of quasi-rent per worker with respect to TFP is:

$$\frac{\partial \ln(\frac{Q_j}{N_j})}{\partial \ln TFP_j} = \frac{VA_j}{Q_j} \times \frac{\partial \ln(\frac{VA_j}{N_j})}{\partial \ln TFP_j} + (w_{Hj}^a - w_{Lj}^a) \frac{w_{Lj}^a L_j + w_{Hj}^a H_j}{Q_j} \times \frac{\partial \ln s_j}{\partial \ln TFP_j}.$$

The first term in this expression can be thought of as giving the (*ceteris paribus*) relative sensitivity of quasi-rents and value added to productivity shocks. Our reading of the literature suggests that the ratio of value added to quasi-rents is around 2, so rent sharing studies that use quasi-rent per worker as the measure of profitability should tend to find elasticities that are about one-half as large as studies that use value added per worker (or a direct measure of TFP). The second term in the expression captures skill upgrading which will tend to augment the relative sensitivity of quasi-rents to productivity shocks in proportion to the gap in alternative wages between type H and L workers. This suggests both that multiplying quasi-rent elasticities by 2 may yield a conservative adjustment and that value added based measures of productivity are less sensitive to neglected worker heterogeneity.

A final approach is to use profits per worker $\frac{\pi_j}{N_j} = \frac{VA_j}{N_j} - w_{Lj}(1 - s_j) - w_{Hj}s_j$ as the rent measure. An equivalent derivation yields:

$$\frac{\partial \ln(\frac{\pi_j}{N_j})}{\partial \ln TFP_j} = \frac{VA_j}{\pi_j} \times \frac{\partial \ln(\frac{VA_j}{N_j})}{\partial \ln TFP_j} + (w_{Hj} - w_{Lj}) \frac{w_{Lj}L_j + w_{Hj}H_j}{\pi_j} \times \frac{\partial \ln s_j}{\partial \ln TFP_j}.$$

Because profits are empirically not much different from quasi-rents, a reasonable adjustment factor is again around 2. As with quasi-rents, estimates based upon profits per worker are more sensitive to neglected worker heterogeneity than value added per worker.

A Summary of the Rent Sharing Literature

Table 1 synthesizes the estimated rent sharing elasticities from the 21 studies listed in Appendix Table 1, 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 based on the level of aggregation in the measures of rents and wages.

The first group of studies, which includes two prominent papers from the early 1990s, uses industry-wide measures of profitability 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, though 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 15 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, though a few studies report rent sharing elasticities that are 0.05 or smaller.

⁵We 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.

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., Card, Cardoso, and Kline, 2016; Carlsson, Messina, and Skans, 2014; Guiso, Pistaferri, and Schivardi, 2005) 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, Devicienti, Maida, 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 microdata, 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.

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 SABI (Sistema de Analisis de Balances Ibericos) 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 data base 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 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 high. Interestingly, the 0.20 to 0.30 range of estimates

⁶Businesses in Portugal are required to file income statements and balance sheet information annually with the Integrated System of Company Accounts. 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, Cardoso and Kline (2016) for more information on the matching process.

⁷A similar finding is reported by Card, Devicienti, and Maida (2014) using Italian data.

is comparable to the range of the studies in the first two panels 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, Pistaferri, and Schivardi, 2005).⁹ The discrepancy is particularly large for OLS models using sales per worker (compare row 2 and row 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, 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 IV estimates in row 6, which point toward a rent sharing elasticity of approximately 0.06.

An interesting feature of both the OLS and IV within-job estimates is that the addition of detailed industry controls reduces the rent sharing elasticity by 10-20 percent. 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 column 3) is somewhat smaller than the elasticity with respect to sector-wide shocks (which are incorporated in the elasticities in the models in column 1). If true more generally, this suggests that the use of industry-wide rent measures will lead to a 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, Messina, and Skans (2014).

Overall, we conclude from the studies in Table 1 and our own within-job estimates for Portugal in Table 2 that a plausible range for the elasticity of wages with respect to value added per worker is 0.05-0.15. Elasticities of this magnitude are clearly too low to rationalize the parallel trends of productivity dispersion and wage dispersion illustrated in Figure 1. When wages contain an employer-specific rent premium, however, wage inequality also depends on the degree of sorting of high- and low-skilled workers to more- and less-profitable employers, which as emphasized in Card, Heining and Kline (2013) can contribute to the trend in wage dispersion.

⁸If 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.

⁹A third potential explanation is selection bias in the stayer models, induced by selecting a sample of job stayers. Results presented in Card, Cardoso and Kline (2016, Appendix Table B10) suggest this factor is relatively small.

2 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 next 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.

AKM Models

In their seminal study of the French labor market, AKM 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 AKM’s 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 striking “symmetry” prediction that we discuss in more detail below.

Estimates of AKM style models on population level administrative datasets from a variety of different countries have found that the firm effects in these models typically explain 15-25 percent 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 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:

$$\begin{aligned} \text{Var}(\ln w_{it}) &= \text{Var}(\alpha_i) + \text{Var}(\psi_{J(i,t)}) + \text{Var}(X'_{it}\beta) + \text{Var}(\varepsilon_{it}) \\ &\quad + 2\text{Cov}(\alpha_i, \psi_{J(i,t)}) + 2\text{Cov}(\alpha_i, X'_{it}\beta) + 2\text{Cov}(\psi_{J(i,t)}, X'_{it}\beta) \end{aligned} \tag{1}$$

which includes both the variance of the firm-specific wage premiums and a term reflecting the covariance

¹⁰For 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 18% and 21% of the variance of the wages of German men depending on the time period studied. Card, Cardoso, and Kline (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 (2015) 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.

of the worker and firm effects. If workers with 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.

An alternative decomposition uses the fact that:

$$Var(\ln w_{it}) = Cov(\ln w_{it}, \alpha_i) + Cov(\ln w_{it}, \psi_{J(i,t)}) + Cov(\ln w_{it}, X'_{it}\beta) + Cov(\ln w_{it}, \varepsilon_{it}). \quad (2)$$

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 $Cov(\ln w_{it}, \psi_{J(i,t)}) = Var(\psi_{J(i,t)}) + Cov(\alpha_i, \psi_{J(i,t)}) + Cov(X'_{it}\beta, \psi_{J(i,t)})$. One way to think about this decomposition is that one half of the firm covariance terms in (1) are attributed to the firm-specific wage premiums.

Identifying Age and Time Effects

A technical issue that arises with the AKM model is appropriate specification of the effects of age (or potential experience). 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, AKM 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 datasets do not always go far enough back to cover the cohorts of interest and because some datasets only report 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, Heining, and Kline (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 (age-40). The same restriction is used in Card, Cardoso, and Kline (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 the results in Card, Cardoso, and Kline (2016) to four alternate normalizations of the age effects. The first column shows the baseline normalization, which attributes a relatively

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

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 (column 2) has little effect on this conclusion, whereas re-normalizing 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 the true profile is everywhere nonlinear. Column 5 shows the results of using a linear combination of normal density functions in age (with five 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 only “solve” the identification problem 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 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 the this literature found under an age 40 normalization are robust to alternate identifying assumptions.

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 the estimated size of $Var(\psi_{J(i,t)})$ or $Var(\alpha_i + X_{it}\beta)$, or the covariance term $Cov(\psi_{J(i,t)}, \alpha_i + X_{it}\beta)$ but, as shown in Table 3, it will affect the estimated covariance of the person and firm effects and the relative size of $Var(\alpha_i)$ versus $Var(\psi_{J(i,t)})$.

Worker-Firm Sorting and Limited Mobility Bias

In their original study, AKM 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, Lengermann, and McKinney (2003) report a correlation of 0.08 for U.S. workers, while Card, Heining and Kline (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

¹²Abowd, Lengerman and McKinney (2003) impose a normalization on the experience profiles in their estimation of an AKM model for the LEHD 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 column 4.

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

Andrews et al. (2012) show convincingly that this “limited-mobility” bias can be substantial. In sampling 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 more robust to unmodeled dependence is an important priority for future research.

Exogenous Mobility

AKM’s 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, Heining, and Kline (2013) (hereafter, CHK) 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, CHK define firm groups based on 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, Cardoso, and Kline (2016). The samples are restricted to workers who switch establishments and have at least two 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 step-like patterns: when workers move to higher paying establishments their wages 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 zero.

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 pre-trend.

¹⁴For 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.

Card, Cardoso, and Kline (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 4 coworker pay groups. Comparisons of upward and downward movers are displayed visually in Figures 5a and 5b, and show that the matched pairs of adjusted wage changes are roughly scattered along a line with slope -1, consistent with the symmetry restriction.

Similar event studies can also be conducted using firm groupings based on the estimated pay effects obtained from an AKM model. As shown in CHK (Figure VII), applying this approach to data for German males yields the same conclusions as an analysis based on coworker pay groups. Macis and Schivardi (2015) conduct this style of event study using Social Security earnings data for Italian workers, and confirm that wage profiles of movers exhibit the same step-like patterns found in Germany and Portugal.

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. Figure 6 and 7, taken from Card, Cardoso and Kline (2016) shows the mean residuals for 100 cells based on deciles of the estimated worker effects and deciles of the estimated firm effects. If the additive model is correct, the residuals should have mean zero 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 zero. Reassuringly, the mean residuals are all relatively close to zero. 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 effects). The largest mean residuals are for the lowest-ability workers in the lowest paying firms – an effect which may reflect the impact of the minimum wage in Portugal. Residual plots for workers and firms in Germany, reported by CHK, and in Italy, reported by Macis and Schivardi (2015), also show no evidence of systematic departures from the predictions of a simple AKM style model.

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.

Though 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 Schmieder (2015) find in large German firms that food, cleaning, security, and logistics (FCSL) 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, Cardoso, and Kline (2016) find that Portuguese women exhibit slightly attenuated firm effects relative to men, which they argue reflects gender differences in bargaining behavior.

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

In their original study AKM 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 higher measured rents hire older workers, or workers with higher 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, Cardoso and Kline (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 (column 1), controls for major industry and city (column 2) and controls for detailed industry and location (column 3). The estimated rent sharing elasticities in row 1 are qualitatively similar to the estimates in row 1 of Table 2 but differ slightly due to differences in the sample arising 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 to 0.15.

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 the gap. One is that the within-job estimates are biased downward

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

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, Pistaferri and Schivardi (2005) is that the structural rent sharing elasticity depends on whether the source of variation in productivity is perceived as permanent or transitory. Pistaferri, Guiso and Schivardi (2005) present a model in which wages adjust less to purely transitory fluctuations than to persistent changes in productivity.¹⁶ To the extent that industry-wide productivity shifts tend to be 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 the extra work effort or less desirable work conditions at higher-productivity firms (e.g., Lavetti and Schmutte, 2015; Sorkin, 2015). If this is true then one would expect the estimated elasticities from row 3 of Table 4 to overstate the true rent sharing effect. Card, Cardoso, and Kline (2016, Appendix 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. Moreover, Card, Heining and Kline (2012, Appendix Table 8) find that jobs with higher wage premiums last significantly longer, suggesting that workers prefer these jobs. Nevertheless, we cannot rule out some role for compensating differentials, suggesting that the estimates in row 3 of Table 4 may overstate the true rent-sharing elasticity.

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 high school or more) in our Portuguese wage sample. We then re-estimated 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 high education group, virtually all of this gap is due to a stronger correlation between the worker quality component of wages and value 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 super-impose a bin-scatter of the relative share of high educated 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

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

¹⁷As shown in Section 1, ignoring variation in capital per worker, the log of value added per worker at firm j is the sum of $\ln(TFP_j)$ and $\ln(q_j)$ where q_j is the productivity-weighted share of higher-skilled workers at the firm, which is increasing in the share of high-education workers. The expected slope of a regression of the log of the relative share of highly educated workers

4 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. Rather than build a model 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, which cannot observe individual worker’s preferences but have knowledge of the population distribution of preferences, with some monopsony power to set wages.¹⁸

While empirical work on monopsony has experienced something of a renaissance (see Manning, 2011 for a review) there has, to our knowledge, 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. 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 (e.g., Burdett and Mortensen, 1998) is that we assume *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

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:

$$v_{iSj} = \beta_S \ln w_{Sj} + a_{Sj} + \epsilon_{iSj},$$

where 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 non-pecuniary match factors such as distance to work or interactions with coworkers and supervisors.²⁰ We assume that the $\{\epsilon_{iSj}\}$ are independent draws

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.

¹⁸In 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.

¹⁹We have also ignored “efficiency wage” explanations for firm wage premia which can emerge, for example, due to 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.

²⁰Note that this specification of preferences is equivalent to assuming that $v_{iSj} = \beta_S^0 \ln w_{Sj} + a_{Sj}^0 + \tau \epsilon_{iSj}$, where $\tau > 0$ is a

from a type I Extreme Value distribution. Given posted wages, workers are free to work at any firm they wish, which leads to logit choice probabilities of the form:

$$p_{Sj} \equiv P(v_{iSj} \geq v_{iSk} \text{ all } k \neq j) = \frac{\exp(\beta_S \ln w_{Sj} + a_{Sj})}{\sum_{k=1}^J \exp(\beta_S \ln w_{Sk} + a_{Sk})}.$$

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:

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

where (λ_H, λ_L) are constants common to all firms in the market. Thus, for large J , the approximate firm-specific supply functions can be written in constant elasticity form as:

$$\ln L_j(w_{Lj}) = \ln(\mathcal{L}\lambda_L) + \beta_L \ln w_{Lj} + a_{Lj} \quad (3)$$

$$\ln H_j(w_{Hj}) = \ln(\mathcal{H}\lambda_H) + \beta_H \ln w_{Hj} + a_{Hj}, \quad (4)$$

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.

Firm optimization

Firms have production functions of the form:

$$Y_j = T_j f(L_j, H_j), \quad (5)$$

where T_j is a firm-specific productivity shifter. We assume that $f(.,.)$ 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 (3) and (4). Firms cannot observe workers' preference draws $\{\epsilon_{iSj}\}$, which prevents them from perfectly price discriminating against workers according to their idiosyncratic reservation values. The firm's optimal wage choices solve the problem:

$$\min_{w_{Lj}, w_{Hj}} w_{Lj} L_j(w_{Lj}) + w_{Hj} H_j(w_{Hj}) \text{ s.t. } T_j f(L_j(w_{Lj}), H_j(w_{Hj})) \geq Y.$$

scale factor, and defining $\beta_S = \beta_S^0/\tau$ and $a_{Sj} = a_{Sj}^0/\tau$.

²¹Berry and Pakes (2007) contrast demand models where consumers have idiosyncratic preferences for specific products versus what they term the "pure characteristics" model where consumers only care about 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 the pure characteristics model is less applicable to the worker's choice of employer because of the many non-pecuniary aspects of work that can give rise to match effects. For example, no two employers have exactly the same location and workplace culture. However, it is clearly an important question for future research which model works better empirically.

²²This 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 $T_j f(L_j, H_j)$. In this case capital and intermediate inputs will adjust proportionally to $T_j f(L_j, H_j)$.

Using (3) and (4), the associated first order conditions can be written as:

$$w_{Lj} = \frac{\beta_L}{1 + \beta_L} T_j f_L \mu_j \quad (6)$$

$$w_{Hj} = \frac{\beta_H}{1 + \beta_H} T_j f_H \mu_j, \quad (7)$$

where μ_j represents the marginal cost of production, which the firm will equate to marginal revenue at an optimal choice for Y . Thus the terms $T_j f_L \mu_j$ and $T_j f_H \mu_j$ on the right hand sides of equations (6) and (7) represent the marginal revenue products of the two skill groups. As noted by Robinson (1933), with upward-sloping supply functions the optimal wages of the two skill groups are *marked down* relative to their marginal revenue products, with “markdown ratios” that depend on their supply elasticities. For example, if $\beta_L = 9$, then the wages of lower-skilled workers will be set at 90% of their 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. The firm’s profits are proportional to the amount by which wages are marked down relative to marginal revenue products. As the elasticities $\beta_L, \beta_H \rightarrow \infty$, the equilibrium approaches the competitive solution and firms choose the numbers of workers in each skill group to equate their marginal revenue products to their market wages.

To proceed, we need to specify the production function and the firm’s marginal revenue function. On the technology side, we start with a simple baseline case where $f(.,.)$ 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 $f(.,.)$ is a CES production function. On the revenue side, we assume that the firm faces a downward-sloping product demand function with elasticity ε , so that marginal revenue is a simple declining function of total output.

Baseline Case: Linear Production Function and Constant Elasticity Demand

To develop intuition, we begin with the simplest possible example where the firm has a linear production function:

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

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. We also assume that the firm’s inverse demand function is $P_j = P_j^0(Y_j)^{-1/\varepsilon}$ with $\varepsilon > 1$ giving the elasticity of product demand. This yields a marginal revenue function:

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

Under this specification of technology and market structure, the first order conditions (6) and (7) evaluate

to:

$$w_{Lj} = \frac{\beta_L}{1 + \beta_L} (1 - \theta) \left(1 - \frac{1}{\varepsilon}\right) T_j P_j^0 Y_j^{-1/\varepsilon} \quad (8)$$

$$w_{Hj} = \frac{\beta_H}{1 + \beta_H} \theta \left(1 - \frac{1}{\varepsilon}\right) T_j P_j^0 Y_j^{-1/\varepsilon}. \quad (9)$$

As a consequence of the linear production technology, the optimal wage choices for the two groups are independent (conditional on marginal revenue and the technology shock). The optimal equilibrium wage in the case of a single skill group is illustrated in Figure 9. On a logarithmic scale, the firm faces an upward-sloping inverse supply function with slope $1/\beta$ (we drop the skill group subscript for simplicity). The associated marginal factor cost (MFC) schedule lies above the inverse supply function but has the same slope.²³ The firm's marginal revenue product (MRP) schedule is downward sloping with constant elasticity $-1/\varepsilon$. The optimal level of employment is determined by the crossing of the MFC and MRP schedules, while the wage is determined from the value of the inverse labor supply function at this level of employment.

Note that the model predicts that more productive firms (i.e., those with higher values of $T_j P_j^0$) will pay higher wages, holding constant the amenity factors a_{Lj} and a_{Hj} . A firm can be more productive either because of firm-specific technological innovations (i.e., shifts in T_j) or firm-specific demand shocks (shifts in P_j^0). Graphically, shocks to either factor will shift the marginal revenue product locus and yield a shift in wages that depends on the relative supply and demand elasticities. It is interesting to derive the resulting rent-sharing elasticities formally. Let $\bar{\beta}_j = \beta_L \kappa_j + \beta_H (1 - \kappa_j)$ denote the average supply elasticity to the firm, where $\kappa_j \equiv \frac{(1-\theta)L_j}{(1-\theta)L_j + \theta H_j}$ gives the share of efficiency units contributed by lower-skilled workers. Then, as detailed in the Appendix, it is straightforward to show that:

$$\begin{aligned} \frac{\partial \ln w_{Lj}}{\partial \ln P_j^0} &= \frac{\partial \ln w_{Hj}}{\partial \ln P_j^0} = \frac{\varepsilon}{\varepsilon + \bar{\beta}_j} \\ \frac{\partial \ln w_{Lj}}{\partial \ln T_j} &= \frac{\partial \ln w_{Hj}}{\partial \ln T_j} = \frac{\varepsilon - 1}{\varepsilon + \bar{\beta}_j}. \end{aligned}$$

Notice that as $\varepsilon \rightarrow \infty$ these elasticities tend to 1 because the marginal revenue product schedule becomes horizontal. On the other hand, as $\bar{\beta}_j \rightarrow \infty$ wages become insensitive to firm-specific demand factors. If $\varepsilon \approx 2$ and $\bar{\beta}_j \approx 9$ we obtain implied elasticities of wages with respect to TFP of roughly 0.09, which is within the range found in the rent sharing literature. Moreover, the predicted equality of the elasticities for high and low skilled workers is consistent with the evidence of equal rent sharing elasticities for low and high education workers in Portugal presented in Table 4.

Changes in the firm-specific amenities for the two groups of workers also affect wages. Visually, an increase in firm-specific amenities leads to an outward shift in the supply and MFC functions, causing a decline in the equilibrium wage that will be larger when the firm's demand is less elastic. More formally:

$$\begin{aligned} \frac{\partial \ln w_{Lj}}{\partial a_{Lj}} &= \frac{\partial \ln w_{Hj}}{\partial a_{Lj}} = \frac{-\kappa_j}{\varepsilon + \bar{\beta}_j} \\ \frac{\partial \ln w_{Lj}}{\partial a_{Hj}} &= \frac{\partial \ln w_{Hj}}{\partial a_{Hj}} = \frac{-(1 - \kappa_j)}{\varepsilon + \bar{\beta}_j}. \end{aligned}$$

²³Denote the labor supply function by $S(w)$, which has elasticity β . The inverse labor supply function is $w = w(S)$ which has elasticity $1/\beta$. Labor cost is $S w(S)$, so marginal labor cost is $w(S)(1 + \beta^{-1})$, which lies above the inverse supply function but has the same elasticity.

Like the effects of T_j and P_j^0 , shifts in group-specific amenities lead to equi-proportional shifts in the wages of the two skill groups. Note that as $\varepsilon \rightarrow \infty$, amenities have no effect on wages because TFP becomes invariant to firm size.

The first order conditions (6) and (7) imply that the *wage gap* between high and low skilled workers is constant across firms:

$$\ln \frac{w_{Hj}}{w_{Lj}} = \ln \frac{\beta_H}{1 + \beta_H} - \ln \frac{\beta_L}{1 + \beta_L} + \ln \frac{\theta}{1 - \theta}. \quad (10)$$

With a fixed relative wage gap the data generating process for individual wages can be written in the form:

$$\ln w_{Si} = \alpha_S + \psi_{J(i)}, \quad (11)$$

where

$$\alpha_S = \ln(1 - \theta) \left(\frac{\beta_L}{1 + \beta_L} \right) \times 1[S = L] + \ln \theta \left(\frac{\beta_H}{1 + \beta_H} \right) \times 1[S = H]$$

is a skill-group-specific constant and

$$\psi_j = \ln \left(\frac{\varepsilon - 1}{\varepsilon} \right) + \ln T_j + \ln P_j^0 - \frac{1}{\varepsilon} \ln Y_j$$

is the firm-specific wage premium paid by firm j . This simple model therefore yields a model for individual wages that is consistent with the additively separable formulation proposed by AKM.

While 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 depends on their relative distribution across firms. In particular:

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

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:

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

The third term in this expression represents a between-firm sorting component of the average wage gap. Card, Cardoso, and Kline (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, Heining and Kline (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.

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 5 different education groups. We normalize the estimated firm effects using the procedure described in Card, Cardoso and Kline (2016), which sets the average firm effect to zero 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, then slightly decreasing.²⁴ Thus, the lifecycle pattern of between-firm sorting contributes to the well-known shape of the lifecycle 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. Thus, a significant share of the wage gap between men with different education levels is attributable to differential sorting.

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 production is in the CES class:

$$f(L_j, H_j) = [(1 - \theta)L_j^\rho + \theta H_j^\rho]^{1/\rho} \quad (12)$$

where $\rho \in (-\infty, 1]$ and the elasticity of substitution is $\sigma = (1 - \rho)^{-1}$. With a CES production function, the marginal productivities of the two groups have the simple form:

$$\begin{aligned} f_L &= (1 - \theta)L_j^{\rho-1}f(L_j, H_j)^{1-\rho} \\ f_H &= \theta H_j^{\rho-1}f(L_j, H_j)^{1-\rho}. \end{aligned}$$

In an abuse of notation we now define $\kappa_j \equiv \frac{(1-\theta)L_j^\rho}{(1-\theta)L_j^\rho + \theta H_j^\rho} = \frac{\partial \ln f}{\partial \ln L_j} = 1 - \frac{\partial \ln f}{\partial \ln H_j}$. As in the case when $f(\cdot, \cdot)$ is linear, κ_j represents the relative contribution of low skilled workers to production.

Assuming that the firm faces a downward-sloping product demand function with elasticity ε , the first order conditions (6) and (7) can be written:

$$\begin{aligned} \left(1 + \frac{1}{\sigma}\beta_L\right) \ln w_{Lj} &= \ln\left(\frac{\beta_L}{1 + \beta_L}\right) + \ln(1 - \theta) - \frac{1}{\sigma}a'_{Lj} + \Gamma_j \\ \left(1 + \frac{1}{\sigma}\beta_H\right) \ln w_{Hj} &= \ln\left(\frac{\beta_H}{1 + \beta_H}\right) + \ln\theta - \frac{1}{\sigma}a'_{Hj} + \Gamma_j \end{aligned}$$

where $a'_{Lj} = a_{Lj} + \ln(\mathcal{L}\lambda_L)$, $a'_{Hj} = a_{Hj} + \ln(\mathcal{H}\lambda_H)$, and $\Gamma_j = \ln(1 - \frac{1}{\varepsilon}) + \ln P_j^0 + (1 - \frac{1}{\varepsilon})\ln T_j + (\frac{1}{\sigma} - \frac{1}{\varepsilon})\ln f(L_j, H_j)$, and we have used the fact that $1 - \rho = \sigma^{-1}$. It follows immediately that if $\beta_H = \beta_L$ then the relative wage ratio $\ln(w_{Hj}/w_{Lj})$ does not depend on P_j^0 or T_j , and demand-side shocks will have the same proportional effects on low- and high-skilled wages.

More generally, some straightforward calculations (shown in the Appendix) establish that:

$$\frac{\partial \ln w_{Lj}}{\partial \ln P_j^0} = \frac{\sigma + \beta_H}{\sigma + \beta_L + \beta_H + (\frac{\sigma}{\varepsilon} - 1)\bar{\beta}_j + \frac{1}{\varepsilon}\beta_L\beta_H} \quad (13)$$

$$\frac{\partial \ln w_{Hj}}{\partial \ln P_j^0} = \frac{\sigma + \beta_L}{\sigma + \beta_L + \beta_H + (\frac{\sigma}{\varepsilon} - 1)\bar{\beta}_j + \frac{1}{\varepsilon}\beta_L\beta_H} \quad (14)$$

where, as above, $\bar{\beta}_j = \kappa_j\beta_L + (1 - \kappa_j)\beta_H$. The corresponding elasticities of wages with respect to physical

²⁴Topel and Ward (1992) showed that job-to-job mobility was an important component of wage growth for young men in the U.S. 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.

productivity (T_j) take the same form but are scaled by $(1 - \frac{1}{\varepsilon})$. Notice that for $\sigma \rightarrow \infty$, we obtain:

$$\frac{\partial \ln w_{Lj}}{\partial \ln P_j^0} = \frac{\partial \ln w_{Hj}}{\partial \ln P_j^0} = \frac{\varepsilon}{\varepsilon + \beta_j},$$

which is the expression derived above for the additive technology case. In fact, when $\beta_L = \beta_H$ equations (13) and (14) simplify to this same expression, independent of the value of σ .

Evaluating the effect of revenue-TFP shocks on relative wages across firms, we have:

$$\frac{\partial \ln(w_{Hj}/w_{Lj})}{\partial \ln P_j^0} = \frac{\beta_L - \beta_H}{\sigma + \beta_L + \beta_H + (\frac{\sigma}{\varepsilon} - 1) \bar{\beta}_j + \frac{1}{\varepsilon} \beta_L \beta_H}. \quad (15)$$

If $\beta_L = \beta_H$ then the wages of the two skill groups are equally responsive to firm-specific product demand shocks or productivity shocks. Otherwise, the wages of the group with the *lower* labor supply elasticity will be more responsive to firm-specific demand shocks. Specifically, if $\beta_H \geq \beta_L$ lower-skilled wages will be *more responsive* to firm-specific variation in productivity than higher-skilled wages. The associated variation in relative employment is:

$$\begin{aligned} \frac{\partial \ln(H_j/L_j)}{\partial \ln P_j^0} &= \frac{\sigma(\beta_H - \beta_L)}{\sigma + \beta_L + \beta_H + (\frac{\sigma}{\varepsilon} - 1) \bar{\beta}_j + \frac{1}{\varepsilon} \beta_L \beta_H} \\ &= -\sigma \frac{\partial \ln(w_{Hj}/w_{Lj})}{\partial \ln P_j^0}. \end{aligned} \quad (16)$$

Thus, variation across firms in relative employment due to firm-specific productivity shocks or product demand shocks is *negatively* correlated with the associated variation in relative wages, with a coefficient reflecting the elasticity of substitution. Equation (16) is similar to the relationship often estimated at the aggregate level between the relative wage of different skill groups and their relative supply.

In the CES case, the relative wages of the two skill groups also depend on the firm-specific amenity shifters a_{Lj} and a_{Hj} . As a point of reference, consider the effect of a parallel shift in the non-wage amenities for both skill groups: $da'_{Lj} = da'_{Hj} = da$. In general, such a shift will lead to a shift in the relative wages of the two groups (with a larger decrease in the wage for the group with the smaller elasticity of supply to the firm). As shown in the Appendix, however, when $\beta_L = \beta_H = \beta$ the associated wage effects are:

$$\frac{\partial \ln w_{Lj}}{\partial a} = \frac{\partial \ln w_{Hj}}{\partial a} = \frac{-1}{\varepsilon + \beta}.$$

Thus, variation in any common amenities will generate proportionate responses in the wages of the two skill groups when their supply elasticities are equal.

As noted in Section 2, existing evidence suggests that the fit of simple AKM-style models is relatively good. In particular, residual plots such as shown in Figures 6 and 7 suggest that the assumption of a constant firm-specific wage premium for different skill groups is not wildly inconsistent with the data for several different European countries. In the context of the models considered here, this suggests either that workers in different skill groups are highly substitutable (i.e., σ is large), or that the supply elasticities of different skill groups are relatively similar and that amenity differences across firms are similar for different skill groups. We suspect that further research along the lines of the simple analysis in Table 4 and Figure 8 will be useful in helping to distinguish these alternative explanations.

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 explain worker flows between firms, which is precisely the variation typically used to estimate firm wage effects. Appendix B 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 as in wage posting models 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 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, 1987; Hwang, Mortensen and Reed, 1998). In our model, non-wage 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 based upon the *elasticity* of labor supply to the firm, which is governed entirely by the distribution of taste heterogeneity. For this reason there are no compensating differentials of the standard sort. Amenities affect firm effects only through their influence on TFP – a firm with attractive non-wage 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 (say) employer provided health benefits on workers than a compensating differentials model.

5 Conclusion

There is no doubt that much of wage inequality is driven by differences in worker skills. But economists have long had evidence (e.g., 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 evidently fail to eliminate productivity and product price differences across firms (Hsieh and Klenow, 2009).

Modern search theory provides one rationale for why firms have some 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' non-wage characteristics. While the mechanisms giving rise to market power under these two approaches are different, both imply that labor is supplied inelastically to firms, which provides them with some scope to set wages. As we have emphasized in our stylized model, the difference between an elasticity of supply to the firm of 9 and an elasticity of infinity turns out to be substantively important for understanding the sources of wage inequality. This is not a difference that can easily be assessed through

introspection, which is why empirical work quantifying the nature of firms' wage setting power is critically important.

The empirical literature on firm wage inequality has progressed dramatically with the introduction of huge matched employer employee datasets. 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 applying (quasi-)experimental research designs that convincingly tease out the mechanisms through which firm shocks are transmitted to workers. 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. A blueprint of sorts is to be found in the industrial organization literature which typically seeks to understand the nature of competition in particular industries, rather than the economy as a whole. It seems plausible that firms have more wage setting power in some labor markets than others and that the nature of firm wage vs. nonwage competition differs as well. How exactly to define a labor market is an important question on which distressingly little work has been done – see Manning and Petrongolo (2011) for one attempt. Nevertheless, some careful case studies of settings where the market structure appears to be clear could be enlightening.²⁵

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.

References

Abowd, J. M., Kramarz, F., & Margolis, D. N. (1999). High wage workers and high wage firms. *Econometrica*, 67(2), 251-333.

Abowd, J., Kramarz, F., Lengermann, P., Perez-Duarte, S. (2004). Are good workers employed by good firms? A test of a simple assortative matching model for France and the United States. Mimeo. CREST.

Abowd, J. M., Lengermann, P. A., & McKinney, K. L. (2003). The Measurement of Human Capital in the US Economy (No. 2002-09). Center for Economic Studies, US Census Bureau.

Abowd, J. A., & Lemieux, T. (1993). The Effects of Product Market Competition on Collective Bargaining Agreements: The Case of Foreign Competition in Canada. *The Quarterly Journal of Economics*, 983-1014.

Acemoglu, D. (2001). Good jobs versus bad jobs. *Journal of Labor Economics*, 19(1), 1-21.

²⁵For example, the empirical literature on monopsony has focused on the market for nurses (e.g., Staiger, Spetz, and Phibbs, 2010) and teachers (e.g., Ransom and Sims, 2010; Falch 2011) based on 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.

- Acemoglu, D., & Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. *Handbook of labor economics*, 4, 1043-1171.
- Akerlof, G. A., & Yellen, J. L. (1986). *Efficiency wage models of the labor market*. Cambridge University Press.
- Albrecht, J. W., & Axell, B. (1984). An equilibrium model of search unemployment. *The Journal of Political Economy*, 824-840.
- Andrews, M. J., Gill, L., Schank, T., & Upward, R. (2008). High wage workers and low wage firms: negative assortative matching or limited mobility bias?. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 171(3), 673-697.
- Andrews, M. J., Gill, L., Schank, T., & Upward, R. (2012). High wage workers match with high wage firms: Clear evidence of the effects of limited mobility bias. *Economics Letters*, 117(3), 824-827.
- Arai, M. (2003). Wages, profits, and capital intensity: Evidence from matched worker-firm data. *Journal of Labor Economics*, 21(3), 593-618.
- Arai, M., & Heyman, F. (2009). *Microdata evidence on rent-sharing*. *Applied Economics*, 41(23), 2965-2976. Chicago
- Ashenfelter, O., & Hannan, T. (1986). Sex discrimination and product market competition: The case of the banking industry. *The Quarterly Journal of Economics*, 149-173.
- Barth, E., Bryson, A., Davis, J. C., & Freeman, R. (2014). It's where you work: Increases in earnings dispersion across establishments and individuals in the US (No. w20447). National Bureau of Economic Research.
- Bagger, J., Christensen, B. J., & Mortensen, D. T. (2014). Wage and labor productivity dispersion: The roles of total factor productivity, labor quality, capital intensity, and rent sharing. Working paper.
- Beaudry, P., Green, D. A., & Sand, B. (2012). Does Industrial Composition Matter for Wages? A test of search and bargaining theory. *Econometrica*, 80(3), 1063-1104.
- Bernard, A. B. & Jensen, J. B. (1995). Exporters, jobs, and wages in US manufacturing: 1976-1987. *Brookings Papers on Economic Activity. Microeconomics*, 67-119.
- Berry, S. (1994). Estimating Discrete-Choice Models of Product Differentiation. *RAND Journal of Economics* 25(2), 242-262.
- Berry, S., & Pakes, A. (2007). The pure characteristics demand model. *International Economic Review*, 48(4), 1193-1225.
- Berry, S., Levinsohn, J., & Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society*, 841-890.
- Blanchflower, D. G., Oswald, A. J., & Sanfey, P. (1996). Wages, profits and rent-sharing. *Quarterly Journal of Economics*, 111(1), 227-251. Chicago
- Black, S. E., & Strahan, P. E. (2001). The division of spoils: rent-sharing and discrimination in a regulated industry. *American Economic Review*, 814-831.
- Boal, W. M. & Ransom, M. R. (1997). Monospony in the the Labor Market. *Journal of Economic Literature* 35(1), 86-112.
- Bonhomme, S., Lamadon, T., & Manresa, E., (2015). A Distributional Framework for Matched Employer Employee Data. Working Paper.
- Burdett, K., & Mortensen, D. T. (1998). Wage differentials, employer size, and unemployment. *International Economic Review*, 257-273.

- Cahuc, P., Postel-Vinay, F., & Robin, J. M. (2006). Wage bargaining with on-the-job search: Theory and evidence. *Econometrica*, 74(2), 323-364.
- Card, D. and Cardoso, A. R. (2012). Can Compulsory Military Service Raise Civilian Wages? Evidence from the Peacetime Draft in Portugal. *American Economic Journal: Applied Economics*, 4(4), 57-93.
- Card, D., Cardoso, A. R., & Kline, P. (2016). Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women. *The Quarterly Journal of Economics*, forthcoming.
- Card, D., Devicienti, F., & Maida, A. (2014). Rent-sharing, Holdup, and Wages: Evidence from Matched Panel Data. *Review of Economic Studies*, 81(1).
- Card, D., Heining, J., & Kline, P. (2012). Workplace heterogeneity and the rise of West German wage inequality (No. w18522). National Bureau of Economic Research.
- Card, D., Heining, J., & Kline, P. (2013). Workplace Heterogeneity and the Rise of West German Wage Inequality. *The Quarterly Journal of Economics*, 128(3), 967-1015.
- Cardoso, A. R. (1997). Workers or employers: who is shaping wage inequality?. *Oxford Bulletin of Economics and Statistics*, 59(4), 523-547.
- Cardoso, A. R. (1999). Firms' wage policies and the rise in labor market inequality: The case of Portugal. *Industrial & labor relations review*, 53(1), 87-102.
- Cardoso, A. R., & Portela, M. (2009). Micro Foundations for Wage Flexibility: Wage Insurance at the Firm Level. *The Scandinavian Journal of Economics*, 111(1), 29-50.
- Carlsson, M., Messina, J., & Skans, O. N. (2014). Wage adjustment and productivity shocks. *The Economic Journal*.
- Charles, K. K., & Guryan, J. (2008). Prejudice and wages: an empirical assessment of Becker's *The Economics of Discrimination*. *Journal of Political Economy*, 116(5), 773-809.
- Charles, K. K., & Guryan, J. (2011). Studying Discrimination: Fundamental Challenges and Recent Progress. *Annu. Rev. Econ.*, 3(1), 479-511.
- Christofides, L. N., & Oswald, A. J. (1992). Real wage determination and rent-sharing in collective bargaining agreements. *The Quarterly Journal of Economics*, 985-1002.
- Ciliberto, F., & Tamer, E. (2009). Market structure and multiple equilibria in airline markets. *Econometrica*, 77(6), 1791-1828.
- Coles, M. G., & Mortensen, D. T. (2016). Equilibrium labor turnover, firm growth and unemployment. *Econometrica*, 84(1), 347-363.
- Dal Bó, E., Finan, F., & Rossi, M. A. (2013). Strengthening state capabilities: The role of financial incentives in the call to public service. *The Quarterly Journal of Economics*, 128(3), 1169-1218.
- Davis, S. J. & Haltiwanger, J. (1991). Wage Dispersion between and within US Manufacturing Plants, 1963-86. *Brookings Papers on Economic Activity. Microeconomics*, 115-200.
- De Menil, G. (1971). *Bargaining: monopoly power versus union power*. MIT Press.
- Du Caju, P., Rycx, F., & Tojerow, I. (2011). Inter-industry Wage Differentials: How Much Does Rent Sharing Matter? *The Manchester School*, 79(4), 691-717.
- Dunne, T., Foster, L., Haltiwanger, J., & Troske, K. R. (2004). Wage and Productivity Dispersion in United States Manufacturing: The Role of Computer Investment. *Journal of Labor Economics*, 22(2), 397-429.
- Eeckhout, J., & Kircher, P. (2011). Identifying sorting—in theory. *The Review of Economic Studies*.
- Estevão, M., & Tevlin, S. (2003). Do firms share their success with workers? The response of wages to product market conditions. *Economica*, 70(280), 597-617.

- Falch, Torberg. (2011). Teacher mobility responses to wage changes: Evidence from a quasi-natural experiment. *American Economic Review* 101 (3), 460-465.
- Faggio, G., Salvanes, K. G., & Van Reenen, J. (2010). The evolution of inequality in productivity and wages: panel data evidence. *Industrial and Corporate Change*, 19(6), 1919-1951.
- Fakhfakh, F., & FitzRoy, F. (2004). Basic wages and firm characteristics: Rent sharing in French manufacturing. *Labour*, 18(4), 615-631.
- Foster, L., Haltiwanger, J., & Syverson, C. (2008). Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?. *The American Economic Review*, 98(1), 394-425.
- Gaure, S. (2014). Correlation bias correction in two-way fixed-effects linear regression. *Stat*, 3(1), 379-390.
- Gibbons, R., & Katz, L. F. (1992). Does unmeasured ability explain inter-industry wage differentials?. *The Review of Economic Studies*, 59(3), 515-535.
- Goldin, C. D., & Katz, L. F. (2009). *The race between education and technology*. Harvard University Press.
- Goldschmidt, D., & Schmieder, J. F. (2015). The Rise of Domestic Outsourcing and the Evolution of the German Wage Structure (No. w21366). National Bureau of Economic Research.
- Groshen, E. L. (1991). Sources of intra-industry wage dispersion: How much do employers matter?. *The Quarterly Journal of Economics*, 869-884.
- Grout, P. A. (1984). Investment and wages in the absence of binding contracts: A Nash bargaining approach. *Econometrica: Journal of the Econometric Society*, 449-460.
- Guiso, L., Pistaferri, L., & Schivardi, F. (2005). Insurance within the Firm. *Journal of Political Economy*, 113(5).
- Gürtzgen, N. (2009). Rent-sharing and Collective Bargaining Coverage: Evidence from Linked Employer–Employee Data. *The Scandinavian Journal of Economics*, 111(2), 323-349.
- Hildreth, A. K., & Oswald, A. J. (1997). Rent-sharing and wages: evidence from company and establishment panels. *Journal of Labor Economics*, 318-337.
- Hildreth, A. K. (1998). Rent-Sharing And Wages: Product Demand Or Technology Driven Premia? *Economics of Innovation and New Technology*, 5(2-4), 199-226.
- Hsieh, C. T., & Klenow, P. J. (2009). Misallocation and Manufacturing TFP in China and India. *The Quarterly Journal of Economics*, 124(4), 1403-1448.
- Hwang, H. S., Mortensen, D. T., & Reed, W. R. (1998). Hedonic wages and labor market search. *Journal of Labor Economics*, 16(4), 815-847.
- Katz, L. F. (1986). Efficiency wage theories: A partial evaluation. In *NBER Macroeconomics Annual 1986, Volume 1* (pp. 235-290). MIT Press.
- Katz, L. F. & Autor, D.H. (1999). Changes in the wage structure and earnings inequality. *Handbook of labor economics*, 3, 1463-1555.
- Katz, L. F. & Summers, L. H. (1989). Industry rents: Evidence and implications. *Brookings Papers on Economic Activity. Microeconomics*, 209-290.
- Krueger, A. B., & Summers, L. H. (1988). Efficiency wages and the inter-industry wage structure. *Econometrica: Journal of the Econometric Society*, 259-293.
- Lavetti K. & Schmutte, I. M. (2015) Estimating Compensating Wage Differentials with Endogenous Job Mobility. Working Paper.

- Lentz, R., & Mortensen, D. T. (2010). Labor market models of worker and firm heterogeneity. *Annu. Rev. Econ.*, 2(1), 577-602.
- Macis, M. & Schivardi, F. (2015). Exports and Wages: Rent Sharing, Workforce Composition, or Returns to Skills? *Journal of Labor Economics*.
- Maré, D. C., & Hyslop, D. (2006). Worker-Firm Heterogeneity and Matching: An analysis using worker and firm fixed effects estimated from LEED. *Statistics New Zealand LEED research paper*.
- Margolis, D. N., & Salvanes, K. G. (2001). Do firms really share rents with their workers?. Working Paper.
- Martins, P. S. (2009). Rent sharing before and after the wage bill. *Applied Economics*, 41(17), 2133-2151.
- Manning, A. (2011). Imperfect competition in the labor market. *Handbook of labor economics*, 4, 973-1041.
- Manning, A., & Petrongolo, B. (2011). How Local Are Labor Markets? Evidence from a Spatial Job Search Model. Working Paper
- Mincer, J. (1974). *Schooling, Experience, and Earnings*. New York: National Bureau of Economic Research.
- Mortensen, D. (2005). *Wage dispersion: why are similar workers paid differently?*. MIT press.
- Murphy, K. M., & Topel, R. H. (1990). Efficiency wages reconsidered: Theory and evidence. In *Advances in the Theory and Measurement of Unemployment* (pp. 204-240). Palgrave Macmillan UK.
- Pyapromdee, S. (2013). *Residual Wage Dispersion with Efficiency Wages*. Working Paper.
- Postel-Vinay, F., & Robin, J. M. (2002). Equilibrium wage dispersion with worker and employer heterogeneity. *Econometrica*, 70(6), 2295-2350.
- Robinson, J. (1933). *The economics of imperfect competition*.
- 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*, University of Chicago Press, 28(2), 331-355.
- Rosen, S. (1987). The theory of equalizing differences. *Handbook of Labor Economics*, 1, 641-692.
- Skans, O. N., Edin, P. A., & Holmlund, B. (2009). Wage dispersion between and within plants: Sweden 1985-2000. In *The structure of wages: an international comparison* (pp. 217-260). University of Chicago Press.
- Slichter, S. H. (1950). Notes on the Structure of Wages. *The Review of Economics and Statistics*, 80-91.
- Song, J., Price, D. J., Guvenen, F., Bloom, N. & von Wachter, T. (2015). Firming Up Inequality (No. w21199). National Bureau of Economic Research.
- Sorkin, I. (2015). *Ranking Firms Using Revealed Preference*. working paper.
- Staiger, D. O., Spetz, J., & Phibbs, C. S. (2010). Is there monopsony in the labor market? Evidence from a natural experiment. *Journal of Labor Economics*, 28(2), 211-236.
- Syverson, C. (2004). Market Structure and Productivity: A Concrete Example. *Journal of Political Economy*, 112(6).
- Syverson, C. (2011). What Determines Productivity?. *Journal of Economic Literature*, 326-365.
- Svejnar, J. (1986). Bargaining power, fear of disagreement, and wage settlements: Theory and evidence from US industry. *Econometrica: Journal of the Econometric Society*, 1055-1078.
- Topel, R. H., & Ward, M. P. (1992). Job Mobility and the Careers of Young Men. *The Quarterly Journal of Economics*, 107(2), 439-479.
- Tirole, J. (1988). *The theory of industrial organization*. MIT press.

Van Reenen, J. (1996). The creation and capture of rents: wages and innovation in a panel of UK companies. *The Quarterly Journal of Economics*, 195-226.

Appendix

A. Derivations

A.1 Linear Technology Substituting $Y_j = f(L_j(w_{L_j}), H_j(w_{H_j}))$ equations (8) and (9) in the text can be written as:

$$\begin{aligned} w_{L_j} &= \frac{\beta_L}{1 + \beta_L} (1 - \theta) T_j^{1-1/\varepsilon} P_j^0 f(L_j(w_{L_j}), H_j(w_{H_j}))^{-1/\varepsilon} \\ w_{H_j} &= \frac{\beta_H}{1 + \beta_H} \theta T_j^{1-1/\varepsilon} P_j^0 f(L_j(w_{L_j}), H_j(w_{H_j}))^{-1/\varepsilon} \end{aligned}$$

To proceed, note that if $f(L_j, H_j) = (1 - \theta)L_j + \theta H_j$ then:

$$\begin{aligned} \frac{\partial \ln f}{\partial \ln L_j} &= \frac{(1 - \theta)L_j}{(1 - \theta)L_j + \theta H_j} = \kappa_j \\ \frac{\partial \ln f}{\partial \ln H} &= \frac{\theta H}{(1 - \theta)L + \theta H} = 1 - \kappa_j \end{aligned}$$

With competitive labor markets, κ_j is the cost share of low-skill labor (i.e., $\kappa_j = \frac{w_{L_j} L_j}{w_{L_j} L_j + w_{H_j} H_j}$). When firms have some market power, however, the cost shares can be different from the relative shares of efficiency units.

Differentiating the first order conditions we obtain

$$\begin{aligned} \begin{bmatrix} 1 + \frac{1}{\varepsilon} \beta_L \kappa_j & \frac{1}{\varepsilon} \beta_H (1 - \kappa_j) \\ \frac{1}{\varepsilon} \beta_L \kappa_j & 1 + \frac{1}{\varepsilon} \beta_H (1 - \kappa_j) \end{bmatrix} \begin{bmatrix} d \ln w_{L_j} \\ d \ln w_{H_j} \end{bmatrix} &= \begin{bmatrix} 1 \\ 1 \end{bmatrix} \left(d \ln P_j^0 + \left(1 - \frac{1}{\varepsilon}\right) d \ln T_j \right) \\ &+ \begin{bmatrix} 1 \\ 1 \end{bmatrix} \left(-\frac{1}{\varepsilon} \kappa_j da_L - \frac{1}{\varepsilon} (1 - \kappa_j) da_H \right) \end{aligned}$$

Define $\bar{\beta}_j = \beta_L \kappa_j + \beta_H (1 - \kappa_j)$ as the average supply elasticity to the firm. Then we obtain:

$$\begin{aligned} \frac{\partial \ln w_{L_j}}{\partial \ln P_j^0} &= \frac{\partial \ln w_{H_j}}{\partial \ln P_j^0} = \frac{\varepsilon}{\varepsilon + \bar{\beta}_j} \\ \frac{\partial \ln w_{L_j}}{\partial a_L} &= \frac{\partial \ln w_{H_j}}{\partial a_L} = \frac{-\kappa_j}{\varepsilon + \bar{\beta}_j} \\ \frac{\partial \ln w_{L_j}}{\partial a_H} &= \frac{\partial \ln w_{H_j}}{\partial a_H} = \frac{-(1 - \kappa_j)}{\varepsilon + \bar{\beta}_j} \end{aligned}$$

The derivatives of wages with respect to the general productivity factor T_j are the same as the derivatives with respect to P_j^0 , but scaled by $(1 - \frac{1}{\varepsilon})$. Notice that as $\varepsilon \rightarrow \infty$ these elasticities tend to 1, as is implied by the assumption of constant marginal revenue.

A.2 CES Technology

We now extend the model by assuming that the firm's production f is in the CES class:

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

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

$$\begin{aligned} f_L &= (1 - \theta)L_j^{\rho-1}f(L_j, H_j)^{1-\rho} \\ f_H &= \theta H_j^{\rho-1}f(L_j, H_j)^{1-\rho} \end{aligned}$$

Define

$$\begin{aligned} \kappa_{Lj} &\equiv \frac{\partial \ln f}{\partial \ln L_j} = \frac{(1 - \theta)L_j^\rho}{(1 - \theta)L_j^\rho + \theta H_j^\rho} \\ \kappa_{Hj} &\equiv \frac{\partial \ln f}{\partial \ln H_j} = \frac{\theta H_j^\rho}{(1 - \theta)L_j^\rho + \theta H_j^\rho} \end{aligned}$$

where $\kappa_{Lj} + \kappa_{Hj} = 1$.

Making use of the above expressions for the marginal products of the two skill groups, and of the marginal revenue function

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

the first order conditions (6) and (7) can be written as:

$$\begin{aligned} \left(1 + \frac{1}{\sigma}\beta_L\right) \ln w_{Lj} &= \ln\left(\frac{\beta_L}{1 + \beta_L}\right) + \ln(1 - \theta) - \frac{1}{\sigma}a'_{Lj} + \Gamma_j \\ \left(1 + \frac{1}{\sigma}\beta_H\right) \ln w_{Hj} &= \ln\left(\frac{\beta_H}{1 + \beta_H}\right) + \ln \theta - \frac{1}{\sigma}a'_{Hj} + \Gamma_j \end{aligned}$$

where $a'_{Lj} = a_{Lj} + \ln(\mathcal{L}\lambda_L)$, $a'_{Hj} = a_{Hj} + \ln(\mathcal{H}\lambda_H)$ and

$$\Gamma_j = \ln\left(1 - \frac{1}{\varepsilon}\right) + \ln P_j^0 + \left(1 - \frac{1}{\varepsilon}\right) \ln T_j + \left(\frac{1}{\sigma} - \frac{1}{\varepsilon}\right) \ln f(L_j, H_j),$$

and we have used the fact that $1 - \rho = \sigma^{-1}$.

The derivatives of the optimal wage choices can therefore be written:

$$\begin{aligned} \begin{bmatrix} 1 + \frac{1}{\sigma}\beta_L - \left(\frac{1}{\sigma} - \frac{1}{\varepsilon}\right)\beta_L\kappa_{Lj} & -\left(\frac{1}{\sigma} - \frac{1}{\varepsilon}\right)\beta_H\kappa_{Hj} \\ -\left(\frac{1}{\sigma} - \frac{1}{\varepsilon}\right)\beta_L\kappa_{Lj} & 1 + \frac{1}{\sigma}\beta_H - \left(\frac{1}{\sigma} - \frac{1}{\varepsilon}\right)\beta_H\kappa_{Hj} \end{bmatrix} \begin{bmatrix} d \ln w_{Lj} \\ d \ln w_{Hj} \end{bmatrix} \\ = \begin{bmatrix} 1 \\ 1 \end{bmatrix} \left(d \ln P_j^0 + \left(1 - \frac{1}{\varepsilon}\right) d \ln T_j \right) \\ - \frac{1}{\sigma} \begin{bmatrix} da_{Lj} \\ da_{Hj} \end{bmatrix} + \begin{bmatrix} 1 \\ 1 \end{bmatrix} \left(\frac{1}{\sigma} - \frac{1}{\varepsilon}\right) (\kappa_{Lj} da_{Lj} + \kappa_{Hj} da_{Hj}) \end{aligned}$$

Some manipulation establishes that:

$$\begin{aligned} \frac{\partial \ln w_{Lj}}{\partial \ln P_j^0} &= \frac{\sigma + \beta_H}{\sigma + \beta_L + \beta_H + \left(\frac{\sigma}{\varepsilon} - 1\right) \bar{\beta}_j + \frac{1}{\varepsilon} \beta_L \beta_H} \\ \frac{\partial \ln w_{Hj}}{\partial \ln P_j^0} &= \frac{\sigma + \beta_L}{\sigma + \beta_L + \beta_H + \left(\frac{\sigma}{\varepsilon} - 1\right) \bar{\beta}_j + \frac{1}{\varepsilon} \beta_L \beta_H} \end{aligned}$$

where, as above, $\bar{\beta}_j = \kappa_{Lj}\beta_L + \kappa_{Hj}\beta_H$. The elasticities with respect to T_j are the same, but scaled by

$(1 - \frac{1}{\epsilon})$. Notice that for $\sigma \rightarrow \infty$, we obtain

$$\frac{\partial \ln w_{Lj}}{\partial \ln P_j^0} = \frac{\partial \ln w_{Hj}}{\partial \ln P_j^0} = \frac{\epsilon}{\epsilon + \bar{\beta}_j},$$

which is the expression derived above for the additive technology case. Moreover, if $\beta_L = \beta_H = \beta$, then

$$\frac{\partial \ln w_{Lj}}{\partial \ln P_j^0} = \frac{\partial \ln w_{Hj}}{\partial \ln P_j^0} = \frac{\epsilon}{\epsilon + \beta},$$

independent of the value of σ .

Finally, consider the effect of a parallel shift in the non-wage amenities for both skill groups: $da'_{Lj} = da'_{Hj} = da$. The effects on the optimal wage choices are:

$$\begin{aligned} \frac{\partial \ln w_{Lj}}{\partial a} &= \frac{-\frac{1}{\epsilon}(\sigma + \beta_H)}{\sigma + \beta_L + \beta_H + (\frac{\sigma}{\epsilon} - 1)\bar{\beta}_j + \frac{1}{\epsilon}\beta_L\beta_H} \\ \frac{\partial \ln w_{Hj}}{\partial a} &= \frac{-\frac{1}{\epsilon}(\sigma + \beta_L)}{\sigma + \beta_L + \beta_H + (\frac{\sigma}{\epsilon} - 1)\bar{\beta}_j + \frac{1}{\epsilon}\beta_L\beta_H}. \end{aligned}$$

A shift in a shared amenity will exert a larger negative effect on the group with the smaller elasticity of supply. In the case where $\beta_L = \beta_H = \beta$, however,

$$\frac{\partial \ln w_{Lj}}{\partial a} = \frac{\partial \ln w_{Hj}}{\partial a} = \frac{-1}{\epsilon + \beta}.$$

B. 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:

$$v_{iSj} = \beta_S \ln w_{Sj} + a_{Sj} + \epsilon_{iSj},$$

where ϵ_{iSj} is drawn from a type I Extreme Value distribution. Hence the period 1 choice probabilities are:

$$\begin{aligned} p_{Sj}^1 &= \frac{\exp(\beta_S \ln w_{Sj}^1 + a_{Sj}^1)}{\sum_{k=1}^J \exp(\beta_S \ln w_{Sk}^1 + a_{Sk}^1)} \\ &\approx \lambda_S^1 \exp(\beta_S \ln w_{Sj}^1 + a_{Sj}^1) \end{aligned}$$

In the second period, a fraction π 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 only choose 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:

$$\begin{aligned} p_{Sj}^2 &= \pi \frac{\exp(\beta_S \ln w_{Sj}^2 + a_{Sj}^2)}{\sum_{l=1}^J \exp(\beta_S \ln w_{Sl}^2 + a_{Sl}^2)} + (1 - \pi) p_{Sj}^1 \\ &\approx \pi \lambda_S^2 \exp(\beta_S \ln w_{Sj}^2 + a_{Sj}^2) + (1 - \pi) p_{Sj}^1 \end{aligned}$$


Note that the period 2 elasticity of supply to the firm is:

$$\frac{\partial \ln p_{Sj}^2}{\partial \ln w_{Sk}^2} = \beta_S \frac{\pi \lambda_S^2 \exp(\beta_S \ln w_{Sk}^2 + a_{Sk}^2)}{\pi \lambda_S^2 \exp(\beta_S \ln w_{Sk}^2 + a_{Sk}^2) + (1 - \pi) p_{Sj}^1}$$

Clearly as $\pi \rightarrow 1$, the elasticity becomes β_S again. Otherwise, we have the elasticity is less than β_S . We also potentially have heterogeneity in the elasticity depending upon how far p_{Sj}^1 is from $\lambda_S^2 \exp(\beta_S \ln w_{Sk}^2 + a_{Sk}^2)$. In a steady state these two objects will be the same and the elasticity of supply to each firm simplifies to $\beta_S \pi$.

Therefore, we can think about the steady state of a dynamic model with taste shocks as being one where firms face a supply curve $\beta_S \pi$ 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 π of workers did in fact draw new tastes. Because the firm cannot budge in its wage policy, each period will yield a fraction π of workers switching between firms.

Figure 1: Trends in Between-Establishment Dispersion in Wages and Productivity



Source: Barth et al (2014)

Figure 2: Implied Age Profiles from AKM Models with Alternative Normalizations of the Age Profile (Men with Primary Education Only)






Figure 3: Mean Log Wages of Portuguese Male Job Changers, Classified by Quartile of Co-Worker Wages at Origin and Destination




Notes: Figure shows mean wages of male workers at mixed-gender firms who changed jobs in 2004-2007 and held the preceding job for 2 or more years, and the new job for 2 or more years. Job is classified into quartiles based on mean log wage of co-workers of both genders. Source: Card, Cardoso and Kline (2016, Figure I).

Figure 4: Mean Wages of Portuguese Female Job Changers, Classified by Quartile of Co-Worker Wages at Origin and Destination




Notes: Figure shows mean wages of female workers at mixed gender firms who changed jobs in 2004-2007 and held the preceding job for 2 or more years, and the new job for 2 or more years. Jobs are classified into quartiles based on mean log wage of co-workers of both genders. Source: Card, Cardoso and Kline (2016, Figure II).

Figure 5a: Test for Symmetry of Regression-Adjusted Wage Changes of Portuguese Male Movers Across Coworker Wage Quartiles




Note: Figure plots regression adjusted mean wage changes over 4 year interval for job changers who move across coworker wage quartile groups indicated. Dashed line represents symmetric changes for upward and downward movers. Source: Card, Cardoso and Kline (2016, Appendix Figure B3).

Figure 5b: Test for Symmetry of Regression-Adjusted Wage Changes of Portuguese Female Movers Across Coworker Wage Quartiles




Note: Figure plots regression adjusted mean wage changes over 4 year interval for job changers who move across coworker wage quartile groups indicated. Dashed line represents symmetric changes for upward and downward movers. Source: Card, Cardoso and Kline (2016, Appendix Figure B4).

Figure 6: Mean Residuals by Person/Firm Deciles,
Portuguese Male Workers




Note: Figure shows mean residuals from estimated AKM model with cells defined by decile of estimated firm effects interacted with decile of estimated person effect. Source: Card, Cardoso and Kline (2016, Appendix Figure B5).

Figure 7: Mean Residuals by Person/Firm Deciles,
Portuguese Female Workers



Note: Figure shows mean residuals from estimated AKM model with cells defined by decile of estimated firm effects interacted with decile of estimated person effect. Source: Card, Cardoso and Kline (2016, Appendix Figure B6).

Figure 8: Relative Wage Premium and Relative Employment of High vs. Low Education Workers



Note: Firms are divided into 100 cells based on mean log value added per worker, 2005-2009, with equal numbers of person-year observations per cell.

Figure 9: Effect of TFP Shock (single skill group)





Figure 10: Mean Firm Effects by Age and Education Group, Portuguese Males



Note: Firm effects are normalized using the method in Card, Cardoso and Kline (2016).

Table 1: Summary of Estimated Rent Sharing Elasticities from the Recent Literature
(Preferred specification, adjusted to TFP basis)

Study and country/industry	Estimated Elasticity	Std. Error
Group 1: Industry-level profit measure		
1. Christofides and Oswald (1992), Canadian manufacturing	0.140	(0.035)
2. Blanchflower, Oswald, Sanfey (1996), US manufacturing	0.060	(0.024)
3. Estevao and Tevlin (2003), US manufacturing	0.290	(0.100)
Group 2: Firm-level profit measure, mean firm wage		
4. Abowd and Lemieux (1993), Canadian manufacturing	0.220	(0.081)
5. Van Reenen (1996), UK manufacturing	0.290	(0.089)
6. Hildreth and Oswald (1997), UK	0.040	(0.010)
7. Hildreth (1998), UK Manufacturing	0.030	(0.010)
8. Barth et al (2014), US	0.160	(0.002)
Group 3: Firm-level profit measure, individual-specific wage		
9. Margolis and Salvanes (2001), French manufacturing	0.062	(0.041)
9. Margolis and Salvanes (2001), Norwegian manufacturing	0.024	(0.006)
10. Arai (2003), Sweden	0.020	(0.004)
11. Guiso, Pistaferri, Schivardi (2005), Italy	0.069	(0.025)
12. Fakhfakh and FitzRoy (2004), French manufacturing	0.120	(0.045)
13. Du Caju, Rycx, Tojerow (2009), Belgium	0.080	(0.010)
14. Martins (2009), Portuguese manufacturing	0.039	(0.021)
15. Guertzgen (2009), Germany	0.048	(0.002)
16. Cardoso and Portela (2009), Portugal	0.092	(0.045)
17. Arai and Hayman (2009), Sweden	0.068	(0.002)
18. Card, Devicienti, Maida (2014), Italy (Veneto region)	0.073	(0.031)
19. Carlsson, Messina, and Skans (2014), Swedish mfg.	0.149	(0.057)
20. Card, Cardoso, Kline (2016), Portugal, between firm	0.156	(0.006)
20. Card, Cardoso, Kline (2016), Portugal, within-job	0.049	(0.007)
21. Bagger et al. (2014), Danish manufacturing	0.090	(0.020)

Note: see Appendix Table 1 for more complete description of each study.

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 2005-2009)			
1. OLS: rent measure = mean log value added per worker 2005-2009	0.270 (0.017)	0.241 (0.015)	0.207 (0.011)
2. OLS: rent measure = mean log sales per worker 2005-2009	0.153 (0.009)	0.171 (0.007)	0.159 (0.004)
3. IV: rent measure = mean log value added per worker 2005-2009. Instrument = mean log sales per worker, 2004-2010	0.327 (0.014)	0.324 (0.011)	0.292 (0.008)
First stage coefficient	0.475 [t=26.19]	0.541 [t=40.72]	0.562 [t=64.38]
B. Within-Job Models (Change in Wages from 2005 to 2009 for stayers)			
4. OLS: rent measure = change in log value added per worker from 2005 to 2009	0.041 (0.006)	0.039 (0.005)	0.034 (0.003)
5. OLS: rent measure = change in log sales per worker from 2005 to 2009	0.015 (0.005)	0.014 (0.004)	0.013 (0.003)
6. IV: rent measure = change in log value added per worker from 2005 to 2009. Instrument = change in log sales per worker, 2004 to 2010	0.061 (0.018)	0.059 (0.017)	0.056 (0.016)
First stage coefficient	0.221 [t=11.82]	0.217 [t=13.98]	0.209 [t=18.63]

Notes: Sample in panel A is 2,503,336 person-year observations from QP for males working in 2005-2009 between the ages of 19 and 65 with at least 2 years of potential experience employed at a firm with complete value added data (from SABI) for 2005 to 2009, and sales data (from QP) for 2004 and 2010. Sample in panel B is 284,071 males age 19-61 in 2005 who worked every year from 2005-2009 at a firm with complete value added data (from SABI) for 2005 to 2009, and sales data (from QP) for 2004 and 2010. Standard errors clustered by firm - 62,845 firms in panel A, 44,661 firms in panel B. Models in panel A control for a cubic in experience and unrestricted education \times year dummies. Models in panel B control for a quadratic in experience and education. Models in column 2 also control for 20 major industries and 2 major cities (Lisbon and Porto). Models in column 3 also control for 202 detailed industry dummies and 29 NUTS-3 location dummies.

Table 3: Summary of Estimated AKM Models for Portuguese Men, Alternative Normalizations of Age Function

	Baseline: Cubic Age Function Flat at Age 40 (1)	Cubic Age Function Flat at Age 50 (2)	Cubic Age Function Flat at Age 30 (3)	Cubic Age Function Flat at Age 0 (4)	Gaussian Basis Function (5)
Std. dev. of person effects (across person-yr obs.)	0.42	0.41	0.46	0.93	0.44
Std. dev. of firm effects (across person-yr obs.)	0.25	0.25	0.25	0.25	0.25
Std. dev. of Xb (across person-yr obs.)	0.07	0.10	0.12	0.74	0.08
Correlation of person/firm effects	0.17	0.16	0.17	0.14	0.17
Correlation of person effects and Xb	0.19	0.19	-0.32	-0.89	-0.06
Correlation of firm effects and Xb	0.11	0.14	-0.03	-0.08	0.04
<i>Inequality decomposition (Percent of variance of log wage explained):</i>					
Person effects + Xb	63	63	63	63	63
Person effects	58	54	70	282	62
Xb	2	3	4	180	2
Cov. of person effects and Xb	3	5	-11	-399	-1
Firm effects	20	20	20	20	20
Cov. of firm effects with (person effect+Xb)	12	12	12	12	12
Cov. of firm effects with person effects	11	10	13	21	12
Cov. of firm effects with Xb	1	2	-1	-9	0
Residual	5	5	5	5	5

Notes: Sample includes 8,225,752 person-year observations for male workers in largest connected set of QP in 2005-2009 period. Sample and baseline specification are the same as in Card, Cardoso and Kline (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. Age function in models in columns 1-4 includes quadratic and cubic terms, with age deviated from 40, 50, 30, and 0 for models in columns 1-4, respectively. Age function in model in column 5 is a Gaussian basis function with 5 equally spaced spline points. All models have the same fit: RMSE of the model is 0.143, the adjusted R-squared is 0.934.

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)			
1. Log Hourly Wage	0.250 (0.018)	0.222 (0.016)	0.187 (0.012)
2. Estimated Person Effect	0.107 (0.010)	0.093 (0.009)	0.074 (0.006)
3. Estimated Firm Effect	0.137 (0.011)	0.123 (0.009)	0.107 (0.008)
4. Estimated Covariate Index	0.001 (0.000)	0.001 (0.000)	0.001 (0.000)
B. Less-Educated Workers (n=1,674,676 person year observations at 36,179 firms)			
5. Log Hourly Wage	0.239 (0.017)	0.211 (0.016)	0.181 (0.011)
6. Estimated Person Effect	0.089 (0.009)	0.072 (0.009)	0.069 (0.005)
7. Estimated Firm Effect	0.144 (0.015)	0.133 (0.013)	0.107 (0.008)
8. Estimated Covariate Index	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)
C. More-Educated Workers (n=577,760 person year observations at 17,615 firms)			
9. Log Hourly Wage	0.275 (0.024)	0.247 (0.020)	0.196 (0.017)
10. Estimated Person Effect	0.137 (0.016)	0.130 (0.013)	0.094 (0.009)
11. Estimated Firm Effect	0.131 (0.012)	0.113 (0.009)	0.099 (0.010)
12. Estimated Covariate Index	-0.001 (0.000)	-0.001 (0.000)	-0.001 (0.000)

Notes: Table 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 clustered by firm in parentheses. Sample in Panel B includes males with less than completed secondary education at firms in the connected set for less educated workers. Sample in Panel C includes males with high school education or more at firms in the connected set for more educated workers. Sample in Panel A includes males in either the Panel B or Panel C sample. All models control for a cubic in experience and unrestricted education \times year dummies. Models in column 2 also control for 20 major industries and 2 major cities (Lisbon and Porto). Models in column 3 also control for 202 detailed industry dummies and 29 NUTS-3 location dummies.

Appendix Table 1: Summary of Estimated Rent-Sharing Elasticities

Study	Design Features	Measure of Profitability	Elasticity
<i>A. Industry-Level Profit Measures</i>			
1. Christofides and Oswald (1992)	Canadian union contracts; 120 narrowly defined manufacturing industries	Industry profits/worker (wage changes)	0.07
2. Blanchflower, Oswald, and Sanfey (1996)	US individual wage data (CPS), grouped to industry×year cells; manufacturing only	Industry profits/worker (within-industry changes)	0.01-0.06
3. 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)	0.29
		Profit per worker (first differences)	0.14
<i>B. Firm-Level Profit Measures, Average Firm-level Wages</i>			
4. 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)	0.22
5. Van Reenen (1996)	Large British manufacturing firms merged with corporate accounts; instruments for rents = innovations, imports, R&D, industry concentration	Quasi-rent/worker (wage change model)	0.29
6. Hildreth and Oswald (1997)	British firms (EXSTAT); firm-specific profits (from financial statements); instruments = lagged values of wages and profits	Profit per worker	0.02
7. Hildreth (1998)	British manufacturing establishments; establishment-specific value added; instruments for rents = innovation measure	Quasi-rent/worker	0.03
8. Barth et al (2014)	US establishments in LBD. Establishment-specific revenues; instrument for revenues/worker = revenues/worker in same industry, other regions	Sales/worker (within-establishment changes)	OLS = 0.32 IV = 0.16

Note: Table continues.

Appendix Table 1 (continued): Summary of Estimated Rent-Sharing Elasticities

Study	Design Features	Measure of Profitability	Elasticity
<i>C. Individual Wages and Firm-Level Profit Measures</i>			
9. 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	France: 0.03 Norway: 0.01
10. Arai (2003)	Swedish worker panel matched to employer (10-year stayers design); profits from financial statements	Change in 5-year average profit per worker	0.01-0.02
11. Guiso, Pistaferri, and Schivardi (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	0.07
		Transitory shock to log value added per worker	0.00
12. Fakhfakh and FitzRoy (2004)	Larger French manufacturing establishments; value added from establishment survey	Mean log value-added/worker over past 3 years	0.12
13. Du Caju, Rycx, and Tojerow (2009)	Belgian establishment panel; value added and labor cost from financial statements	Value added minus labor costs per worker	0.03-0.04
14. Martins (2009)	Larger Portuguese manufacturing firms; revenue and capital costs from financial statements; instruments =export share of sales × exchange rate changes	Revenue-capital costs/worker (differenced)	0.03-0.05
15. Guertzen (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)	0.03-0.04
		Change in quasi-rent/worker (stayers design)	0.01-0.06

Note: Table continues.

Appendix Table 1 (continued): Summary of Estimated Rent-Sharing Elasticities

Study	Design Features	Measure of Profitability	Elasticity
<i>C. Individual Wages and Firm-Level Profit Measures (continued)</i>			
16. Cardoso and Portela (2009)	Portuguese worker panel; sales from firm reports; model-based decomposition of sales shocks	Permanent shock to log sales	0.09
		Transitory shock to log sales	0.00
17. Arai and Hayman (2009)	Swedish worker/firm panel; profits from financial statements; stayers design; instrument =change in foreign sales	Change in profit per worker	0.07
18. Card, Devicienti, and Maida (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)	0.06-0.08
19. Carlsson, Messina, and Skans (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)	0.05
		Sectoral average output/worker (within-job-match)	0.15
20. Card, Cardoso, and Kline (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	Males: 0.16 Females: 0.14
		Mean value added per worker (changes for stayers)	Males: 0.05 Females: 0.04
21. Bagger, Christensen, and Mortensen (2014)	Danish worker panel matched to firms; output from firm survey; non-parametric regressions within sector of wages on labor productivity	Output per worker	Manuf: 0.09 Trade: 0.13 Transp/Comm: 0.05 Finance/Real Est: 0.07

Notes: estimates extracted by authors from studies listed.