

Economics Synopsis

The sources of wage variability in Portugal: a binge reading survey

Pedro Portugal
Banco de Portugal e NOVA-SBE

April 2020

Abstract

The availability of repeated observations on the wages of the workers, across different firms over a long period of time, enable the investigator to explore relevant information about *who the worker is* and *where he works*. When a worker establishes a match with a firm his wage will be partially determined by who he is, where he works, and by the success of the idiosyncratic match. This third component is called *match quality*. Besides the notion of who is the worker, where he works, and the quality of the match, arguably, there is an important dimension often neglected in the literature: *what he does*. In this survey, nine empirical studies exploring the sources of wage variability in the Portuguese labor market are summarized. The common feature of these studies is the attention given to worker, firm, and job-title heterogeneity.

1. The brave new world of worker, firm, match, and job title heterogeneity

In a perfectly competitive labor market, identical workers, working in homogeneous firms, should receive exactly the same wage, which would mean that the wages distribution would collapse to a single mass point. Even in the presence of search costs, the equilibrium outcome would still be unique, corresponding to the monopsony wage, a result known as the Diamond paradox, after Peter Diamond's (1971) celebrated paper. Observed wages, however, reveal an unsuspected degree of wage dispersion, even among observationally identical workers.

The theoretical construct of an equilibrium wage dispersion can, of course, be rationalized in a number of ways by considering heterogeneity across workers and/or firms either ex-ante or as the outcome of a monopsonistic wage competition (most notably under on-the job search, as in Burdett and Mortensen, 1998).

Acknowledgements: I am deeply thankful to Lucena Vieira for so patiently assisting me exploring the Quadros de Pessoal over the last 20 years. I am also grateful to António Antunes and Nuno Alves for valuable comments and suggestions. The analyses, opinions, and findings expressed in this article are those of the author and do not necessarily coincide with those of Banco de Portugal or the Eurosystem.

Note: This article was submitted for approval on February 21, 2020, before the World Health Organization declared the Covid-19 a pandemic on March 11.

E-mail: pportugal@bportugal.pt

The empirical contribution of worker and firm heterogeneity to wage formation has been studied extensively, in particular after the seminal work of Abowd, Kramarz, and Margolis (AKM) (1999). By extending the conventional Mincerian wage equation to include worker and firm fixed effects, the study of AKM spurred a significant interest on the role of firms setting wages, uncovering the importance of heterogeneity of the firms' wage setting policies.

The access to matched employer-employee datasets, by enabling the researcher to follow the worker in the firm and his mobility over firms, allows the econometrician to measure and control for permanent worker and firm heterogeneity. In other words, by having repeated observations on the wages of the workers across different firms over a long period of time, the investigator grabs relevant information about who the worker is and where he works. In AKM models, this information is summarized in the worker and firm fixed effects. Worker effects are, of course, no more than averages of residuals, exhibiting whether or not the worker collects wages above or below what would be predicted by his observed characteristics. Similarly, firm fixed effects reflect the (assumed) constant generosity of the firms' wage policies.

What lies behind a worker fixed effect? Whatever respects the individual worker, is constant over the sample period, and influences the level of wages. Observed worker characteristics which do not change during the survey period such as gender, race, or birth cohort are implicitly accounted for in the wage regression with worker fixed effects. But permanent unobserved characteristics are also accounted for. Notoriously, the role of the workers' ability is subsumed in the worker fixed effect as is his family background or his motivation. Other features may also be unobserved to the econometrician but play a role driving the worker wage. For example, there is no shortage of studies showing that height, weight, or beauty have an effect on wages. The same can be said with respect to religion preferences or sexual orientation. The worker fixed effect catches all of those influences.

What lies behind a firm fixed effect? Whatever respects the firms' wage policy that persists over the survey period. Observed characteristics such as location and sector of activity are likely to affect wages and remain unchanged. Managerial ability, however, may well be unobserved. Economic rents, total factor productivity, and monopsony power may have an impact on wages but be unknown to the econometrician. Again, the firm fixed effect will condense all those effects.

Besides the notion of who is the worker and where he works, arguably, there is an important third dimension often neglected in the literature: what he does. This means that occupational or job title heterogeneity should also be considered in a wage regression, possibly by adding a third high-dimensional fixed effect. The presence of compensating wage differentials associated with the amenities or the hardship of the workplace would justify the inclusion of a job title fixed effect to the extent that it affects persistently the level of wages. Similarly, the risk of accidents or the health risks associated with specific occupations are also likely to engender compensating differentials. Union bargaining power may also be manifested at the job title level because collective agreements over wage floors are negotiated precisely at this level. Finally, the chronic overcrowding in some occupations will be translated into

persistently lower wages for those occupations. The job title fixed effect is expected to absorb all those observed and unobserved impacts.

Another useful extension of this framework, one that relies on the estimation of worker fixed effects and the identification of their co-workers, is to consider the human capital spillovers that may emerge from the relationship with other workers at the workplace. In this vein, the researcher wants to consider not only who the worker is, where he works, and what he does, but also with whom. The presence of peer effects may arrive from learning externalities or simply from peer pressure. In any case, the inclusion of the worker fixed effect in this extended wage regression circumvents the problem of homophily (workers sharing similar characteristics) and the inclusion of firm and job title fixed effects bypasses the problem of spurious correlation at the workplace level (Mansky, 1993; Angrist, 2014).

The allocation of workers to distinct firms and job title is not random. Indeed, workers search for employers and choose occupations maximizing the expected future stream of labor earnings (Roy, 1951). The inclusion of firm and job title fixed effects in the wage regression allows the econometrician to control for the endogenous sorting across firms and job titles, to measure the importance of sorting, and to construct counterfactual exercises where workers are randomly assigned to firms and/or job titles.

The estimation of worker and firm fixed effects enables the researcher to investigate how high and low wage workers match with high and low wage firms. The correlation coefficient between the worker and the firm fixed effects, if positive, will provide evidence that low wage workers are matched with low wage firms and that high wage workers populate disproportionately high wage firms. Until recently, the evidence on the sign of wage sorting was blurred, not least because measurement error in one fixed effect is transmitted with the opposite sign to the other fixed effect, biasing negatively the correlation coefficient.

When a worker establishes a match with a firm his wage will be partially determined by his fixed effect (who he is), by the firm fixed effect (where he works) and by the success (or failure) of the idiosyncratic match. To this third component we shall call match quality. In general, the match quality component cannot be disentangled from the work and firm components. However, under the assumption of orthogonality between the match quality and the worker and the firm fixed effects, a lower bound estimate of the match quality can be computed (Woodcock, 2015). In some instances, considering the match quality component can be very revealing.

In principle, the indication of a positive correlation between worker and firm fixed effects should not be confused with the notion of positive assortative matching, where high productivity workers are assigned to high productivity firms. Whereas it may be argued that the worker fixed effect reflects to a large extent his underlying productivity, there is no reason to establish that high productivity firms are necessarily high paying firms. In the absence of direct measure of firm's productivity, the sign and magnitude of assortative matching would be hard to grasp.

The inclusion of worker, firm, and job title fixed effects in the wage regression equation enlarges dramatically the set of research questions that can be addressed. In particular, the extended model allows the researcher to investigate the role of observed

and, in particular, the unobserved characteristics of the worker, the firm, and the job title. A useful methodology to study the impact of enclosing (persistent) unobserved heterogeneity is the extension suggested by Cardoso et al. (2016) of the decomposition method proposed by Gelbach (2016) in order to account for high-dimensional fixed effects. With this methodology one can ask, for example, by how much the gender wage gap would be reduced if workers were randomly assigned to firms? Is it true that the union wage premium is driven by the fact that unionized workers are more productive? Which part of the wage loss of displaced workers is due to the loss of match quality in the pre-displacement job? Do low productivity workers lose jobs disproportionately during recessions? Do workers move from high paying job titles to low paying job titles during downturns? How much of the return to schooling arises from working with more educated workers?

In this survey, nine empirical studies applied to the Portuguese labor market will be considered, where the wage equation including the three (high-dimensional) sources of heterogeneity (worker, firm, and job-title) is estimated and where the questions listed above, and others, are addressed.¹

2. The architecture of the Portuguese wage setting system

The Portuguese Constitution gives registered trade unions the exclusive right to bargain over wages. Collective bargaining occurs mainly at the industry level but, irrespective of the level, collective agreements on wages determine wage floors only. According to the Labor Code, these union negotiated wage agreements only formally bind the members of employers association and the workers of signatory trade union(s). Collective agreements are of three types. First, and most importantly, there are industry level agreements negotiated between one or more employers' associations and one or more unions, known as *Contratos Coletivos de Trabalho* (CCTs). Second, there are collective agreements signed by several employers that are not part of an employers' association and one or more trade unions, known as *Acordos Coletivos de Trabalho* (or ACTs), that are significant in the financial sector and utilities. Finally, there are firm-level agreements between an individual company and one or more unions, so-called *Acordos de Empresa* (or AEs).

More concretely, collective wage agreements specify wage floors for the base wage for a set of job titles (*categorias profissionais*). Job titles are occupational categories whose specificity is sufficient to justify to bargain over. The detailed classification of the job titles accounts for the complexity of the tasks, the hierarchical standing of the worker, and the stress of the working conditions. In a typical year, around 300 collective agreements are signed covering close to 30000 job titles. Note that this classification is much more granular than commonly used occupation disaggregation.

1. The computational difficulties surrounding the estimation of this type of models are discussed in Guimarães and Portugal (2010), in which an algorithm overcoming those difficulties is offered.

Despite the legal strictures, the impact of collective bargaining agreements reaches far beyond the signatory parties. The critical mechanism shaping the formation of wages has been the systematic extension, via so-called *Portarias de Extensão*, of industry-wide agreements by the Ministry of Employment, following a request from either or both of the parties to the agreement. This means that, at the firm level, there are no longer wage differences between union and non-union members. Similarly, there will be no distinction between members and non-members of the employer association in terms of the wage agreement. The upshot of this near automatic procedure is that the main channel of union strength (as indexed by density) can only be observed at the collective agreement (typically, at the industry) level. Altogether, both agreements and their subsequent extensions explain levels of collective bargaining coverage in the order of 90 percent of workers. Firms frequently pay more than the bargained wage. In practice, around 40 percent of workers receive base wages above the bargained wage. Moreover, on top of the base wage, workers may receive meal allowances, overtime compensation, shift pay, and productivity and non-productivity related bonuses.

The difference between total wages and bargained wages is commonly termed *wage cushion*. There is some evidence that the wage cushion is more relevant in low paying sectors (see, in particular, Cardoso and Portugal, 2005). The explanation seems to be that employers in low bargained wage settings (most likely, industries with low union bargaining power) have the ability to set wages above the minimum to suit their particular labor market conditions.

Apart from this disaggregated or informal minimum wage apparatus, the share of workers receiving the national minimum wage has risen dramatically in recent years (the incidence of minimum wage earners currently exceeds 20 percent). Collective agreements and therefore extension arrangements can only pay less than the national minimum wage to specified types of workers such as apprentices and handicapped persons.²

3. A unique matched employer/employee/collective agreement/job title dataset

Researchers of the Portuguese labor market have been blessed with the access to a uniquely rich dataset which, by construction, matches workers with establishments while registering the worker's job title and the corresponding collective wage agreement.

The Portuguese data employed in the surveyed studies come from a longitudinal matched employer-employee dataset known as the Tables of Personnel (or *Quadros de Pessoal*) for the years 1986 to 2009 (excepting 1990 and 2001) and from its virtually identical successor survey the Single Report (or *Relatório Único*) for the years 2010 to 2017. This unique dataset is administered by the Portuguese Ministry of Employment,

2. For a detailed review of recent developments in collective bargaining and extension arrangements in Portugal, see Addison et al., 2017.

and is taken from a mandatory annual survey of all firms with at least one wage earner in the reference month, March of each year until 1993, October thereafter. The survey covers various firm and establishment characteristics, as well as a set of characteristics of the workforce. Being compulsory, it does not suffer from the non-response problems that often plague standard household and firm surveys. Further, the survey covers all Portuguese wage earners, with the exceptions of the Public Administration sector and domestic servants.

More specifically, the dataset includes information on the establishment (identifier, location, industry, and employment), the firm (firm identifier, location, industry, legal form, ownership, year of formation, employment, sales, and capital), and its workers (social security identifier, gender, age, education, skills, occupation, employment status, professional level, seniority, earnings [base wage, seniority-related earnings, other regular and irregular benefits, and overtime pay], normal and overtime hours, time elapsed since last promotion, professional category and the corresponding classification in a collective agreement).

4. The sources of wage variation

Why do similar workers receive different remuneration and why do similar firms pay different wages? In a labor market operating under perfect competition, each worker should receive a wage that equals his or her marginal (revenue) product. Wage differentials should reflect differences in worker productivity rather than depend on job or employer attributes (other than those affecting worker utility such as dangerous working conditions that will in normal circumstances attract a compensating differential). In turn, worker productivity has a basis in competence, whether observed or not, typically 'acquired' through investments in human capital.

There is no shortage of models seeking additional or alternative explanations for wage variability, but it is now the characteristics of firms that assume pole position. Given the plethora of such treatments (e.g., implicit contract theory, principal-agent models, and efficiency wages), two approaches take front stage. The first has a basis in rent-sharing/insider-outsider considerations, while the second emphasises labor market frictions.

Rent-sharing models predict that wages depend on the employer's ability to pay. In particular, wages are predicted to have a positive correlation with firm profits, since firms may find it beneficial to share their gains with their workers. These models explain why wages depend not only on external labor market conditions but also on firm productivity, profits, degree of competition, and turnover costs, as well as the bargaining strength of workers. They also explain why the wages of workers from different groups of occupations, educational categories, and seniority tiers are higher in some firms or industries than in others.

The other explanation for wage differentials among workers with similar characteristics targeted here derives from the job search and matching literature and emphasises the role of labor market frictions in wage determination. Thus, the

equilibrium job search model of Burdett and Mortensen (1998) predicts that firms may have incentives to offer higher wages than their competitors in order to guarantee a low quit rate and attract a large number of workers in a market characterized by the existence of frictions – even in circumstances of homogeneous workers and firms *ex ante*. For their part, matching models that also take into account the existence of frictions in the labor market provide an explanation for wage dispersion. Differences in match productivity in those models explain gains with their workers and pay above the going rate.

Another source of wage variation is job title heterogeneity, which may influence wage rates for a variety of reasons. First, it is well-known that tasks involving risks of fatal or otherwise serious accidents are better paid than safe tasks. Thus, for example, one should expect significant compensating differentials attached to occupations such as bullfighters (included in our sample). Second, jobs that need to be executed under difficult or stressful conditions are also expected to be more highly remunerated than jobs performed in pleasant environments; for example, one should observe higher wages for individuals working on offshore oil platforms or in mines. Third, the complexity of some tasks may require heavy doses of specific training or unusually skilled workers. This would be a reason why, for example, brain surgeons and prima ballerinas earn higher wages. Fourth, some occupations are known to be chronically ‘overcrowded’, whereas others are thought to be in excess demand. Finally, the kind of technology in use may also foster unionization of the workplace and favor rent seeking. Production activities that imply the concentration of a large number of workers in a single plant (say, in auto or shipbuilding industries) facilitate industrial action, and thus improved wage conditions. To properly incorporate these and other such wage determinants, one needs a very detailed accounting of the kind of jobs being undertaken by workers.

The work by Torres et al. (2016) were the first to explore the relative importance of worker, firm, and job title heterogeneity. The authors reported that worker heterogeneity is the most important source of wage variation in Portugal (contributing to one third of the wage variation). The unobserved component plays a slightly less important role than the observed non-time varying characteristics of workers such as gender and education. Firm effects were also found to be important (contributing one fourth), due in roughly equal parts to the unobserved component and to observed non-time varying characteristics such as regional location, capital ownership, and industry. For their part, although less important than either of the corresponding worker or firm effects, job title fixed effects still explain an unsuspectingly important one fifth of wage variation.

5. Wage sorting in the Portuguese labor market

In the Portuguese labor market, there is convincing evidence that the high wage workers sort themselves into high wage firms (Torres et al., 2016). The evolution of wage sorting, however, is quite striking. Torres and co-authors show that the correlation between the worker and the firm fixed effect steadily declined from 0.31 in 1991 to 0.08 in 2013. This evolution is in sharp contrast with the evidence presented for Germany by Card et al. (2013), which point to a sharp increase in wage sorting over time.

A number of labor market trends may have contributed to the decrease of wage sorting over the last 20 years of the survey. First, the average size of Portuguese firms has declined noticeably; second, the fraction of minimum wage earners has increased noticeably; third, the proportion of employed women has increased; and, finally, the fraction of college graduates has increased markedly. However, when each of the compositional change is isolated, there is no indication that the reported trends lie behind the evolution of the correlation coefficients. Indeed, it is apparent that the decrease of wage sorting is observed within each category. In particular, the trend towards less positive wage sorting occurs primarily among high wage workers.

6. Assortative matching in the Portuguese labor market

The sorting of heterogeneous workers across firms has been the subject of not inconsiderable debate. The idea behind positive assortative matching is the complementarity between individual and plant productivity levels, with good workers being teamed up with good firms. The theoretical basis for such matching is provided by assignment models. The early assignment models, however, were rooted in competitive equilibrium, thereby disregarding establishment-specific components in the wage equation. With the introduction of frictions, more recent developments have ensured a sorting of workers across plants. At issue in these models is the nature of the equilibrium matching pattern since different matching models predict different patterns according to the assumptions of the model.

There is a general consensus that good workers (i.e. the more productive ones) tend to earn higher wages. Therefore, it is possible to rank workers' productivity based on the individual permanent component of their wages, namely the worker fixed effects estimated from wage equations. Similarly, good firms (i.e. more productive ones) tend to have higher profits. However, these firms may pay lower or higher wages due to the presence of non-monotonicities in the wage schedule. Indeed, high-productivity firms have better outside options than their low-productivity counterparts, which may exert downward pressure on their workers' wages. This can be particularly relevant for low-skilled workers who may end up being paid less than if they were working for less productive firms. Non-monotonicities in the wage schedule also mean that wages reflect the marginal contribution to the value that the firm generates and it can be either the more productive or the less productive firms that derive higher marginal benefit from employing a better worker (Eeckhout and Kircher, 2011).

The estimation results given in Torres et al. (2018), in which the labor elasticities of a production function with thousands of different labor inputs are estimated, conveyed a consistent story in favor of the super-modularity or positive assortative matching hypothesis. The similarity between the magnitude of the wage sorting and the assortative matching correlations and, revealingly, their evolution over time seem to suggest that the correlations between the estimates of person and firm fixed effects from the wage equation may after all provide a sensible approximation to the measure of assortative matching. Contrary evidence from similar studies may be tied to the short

temporal dimension of the panels used. As argued earlier, fixed sample and limited mobility biases seem the likely culprits.

7. The sources of wage dispersion

The estimation of a wage regression with high-dimensional fixed effects provides a natural framework to investigate the sources of wage dispersion, and its evolution through time. The decomposition of the elements of the variance of wages is informative regarding the forces behind the changes in wage inequality.

When the breakdown of variance of (log) wages on an annual basis is established, Portugal, et al. (2018) show that the contribution of the main components in the variance of wages is essentially constant over the period under review. Indeed, it is not possible to establish a clear trend for wage dispersion when we consider each wage component separately (the variances of the worker, firm, and job title fixed effects).

It is not surprising, then, that the conventional indicators of wage dispersion have remained constant for the past 20 years. As mentioned above, the correlation between the fixed effect of the worker and the fixed effect of the firm, however, has been weakening, which in turn has led to less wage inequality.

8. The sources of real wage cyclicalities

The cyclical behavior of real wages has been the subject of many studies since the debate of Keynes, Dunlop, and Tarshis. Earlier studies based on aggregate data showed some ambiguous results. One reason why these studies have reached no definitive conclusions resides in the fact that they have ignored the changes in the composition of the workforce over the cycle. Furthermore, with aggregate data it is implicitly assumed that the relationship between real wages and the business cycle is the same for all individuals or groups of individuals. If this assumption does not hold, then the estimates of real wage cyclicalities are plagued by a specification bias.

The presence of compositional effects has attracted much attention in recent years and the seminal studies based on individual-level panel data for the United States show that cyclical changes in the composition of the workforce may generate a bias in the aggregate real wage.

It is also widely agreed that industry composition may change over the cycle. As pointed out by Okun, if some industries/firms offer rents to workers and if these sectors are also more cyclically sensitive, workers can switch into high-paying jobs during booms because such jobs are less tightly rationed during these times. This inter-industry/firm mobility of workers generates, per se, a procyclical behavior of wages.

After controlling for worker and firm heterogeneity, it can always be argued that the composition and quality of jobs within a firm are also likely to vary over the business cycle. If firms' promotions and hiring standards exhibit a cyclical pattern, overall wage cyclicalities is mainly driven by workers changing job titles rather than by wage changes within job titles. The same line of reasoning applies if match quality among new hires

falls in a recession or firms hire proportionally more workers into low-skilled jobs in a recession than in an expansion.

In the study of Carneiro et al. (2012) it is shown that, accounting simultaneously for worker, firm, and job title heterogeneity, the data support the hypothesis that real wages in Portugal are quite sensitive to the business cycle, regardless of the measure of the cycle used. Compositional bias, however, plays a very important role. Failure to account for worker, firm, and job title unobserved heterogeneity induces a countercyclical bias in wage cyclicalities. This is due to the fact that low-skilled workers and low-productivity/low-paying jobs account for a smaller share of employment in recessions than in expansions.

An important finding presented by Carneiro et al. (2012) was that entry wages are more responsive to the cycle than wages of existing workers within the same firm-job title. This finding is valid whether we use the unemployment rate as the measure of the business cycle fluctuations or the most common measure in search and matching models, the aggregate labor productivity. It is also shown that, after account for worker, firm, and job title heterogeneity, wages for all types of workers exhibit a wage-productivity elasticity close to unity. Finally, disentangling the job finding and job separation probabilities, real wages react positively to changes in the job finding probability and negatively to changes in the job separation probability.

9. The sources of gender discrimination

In a market where there are good and bad firms, or simply firms with different wage policies, segregation of workers across firms is likely to lead to a wage gap that will persist over time (Groschen, 1991).

Similarly, the allocation of workers to differently payed occupations or job titles, in particular, the restrictions to access those that are highly payed (glass ceilings), may also generate wage penalties.

Cardoso et al. (2016a) investigate the mechanisms that shape the gender wage gap in Portugal and provide a clear measure of the impact of the allocation of workers to firms and job titles. They found that one-fifth of the gender gap results from the segregation of workers across firms, and another one-fifth results from job title segregation.

Not only are women sorted more frequently into low-wage firms, but also the wage penalty increases with the size of the firm fixed effect. The authors conclude that the 'glass ceiling effect' operates mainly through worker allocation to firms rather than occupations.

Why are women penalised by allocation across firms? One possibility would be that female workers are less efficient job searchers. A number of mechanisms may lie behind this. First, women may search less intensively than men because they may have more productive alternatives for the allocation of their time (e.g., domestic production). Second, women may limit the set of job opportunities to jobs with particular characteristics (e.g., flexibility of work schedules). Third, women may shape their search strategy to the labour supply decisions of their husbands, as implied by

first generation labour supply models. For example, those women may severely restrict the geographical boundaries of their job search. Fourth, women may have significant disadvantages compared with men, exploiting their social networks to locate suitable job offers. Fifth, women may underestimate the relevant distribution of wage offers, either because they expect to be discriminated against by some firms, or because they misrepresent their true value to the firm and bargain wages less aggressively than men (Card et al., 2015). A telling indication that women have lower expectations regarding wages is given by evidence that they report lower reservation wages than men for similar observed characteristics (Addison et al., 2004).

Why are women allocated to lower-paying job titles? What is behind job title segregation? The presence of barriers to entry into high-paying job titles, driven either by the hiring decisions of the employers or by the requirements of the job, is certainly one of the mechanisms at work. Job promotion decisions biased against female workers may also be at play. Overcrowding of job titles highly preferred by female workers may drive their corresponding wages down. On the other hand, the access to some high-wage job titles that are controlled by closed shop trade unions are frequently male-dominated (e.g., longshoremen). The allocation into job titles is responsible for around one-fifth of the total gender gap. There is, however, no indication that the wage gap between men and women increases significantly along the job title fixed effect dimension.

Of concern is the indication of the study by Cardoso et al. (2016b) warning that the allocation of female workers to firms and job-titles did not improve over the last two decades. In fact it deteriorated somewhat. By 2013, females tend to be less present in firms and job titles with more generous wage policies, in comparison with 1991.

10. Monopsony power in the Portuguese labor market

The assumption of a single market wage that would cause all employees to instantaneously leave the firm after a one cent wage cut seems unrealistic. The empirical evidence suggests the presence of considerable wage dispersion among workers with similar characteristics and among similar firms. In particular, the firm effects estimated in wage regressions can be thought of as arising from distortions in the labor markets. Search frictions in the labor market such as imperfect information on alternative available jobs, moving and learning costs, firm specific human capital, reputation costs, exploitation of rents, and worker heterogeneous preferences over non-wage job characteristics, are sources of labor market power, and help explain why firms have market power and why the labor supply curve faced by an individual firm is not perfectly elastic. These search frictions in the labor market may generate upward sloping labor supply curves to a particular firm.

In a standard wage setting model this means that firms possess some power to mark down their wages below the marginal revenue product. This is in line with the “new monopsony” literature popularized by Manning (2003), in which employers gain some market power derived from search frictions when setting wages. Monopsony is not understood in the traditional sense of a unique employer in the labor market,

but instead as synonymous with imperfect competition, monopsonistic competition, upward sloping labor supply curve to the firm, or finite labor supply elasticity. A particular firm may face an upward labor supply curve even if there is no concentration on the demand side of the market.

Félix and Portugal (2017) estimate a measure of labor supply elasticity to the firm directly from the production function, at a very granular level. Investigating the impact of the elasticity of labor supply to a particular firm and firm total factor productivity on individuals' earnings, the authors use the Gelbach exact decomposition to understand how firm's monopsony power is associated with the firm's wage setting policy. The results suggest that a one standard deviation increase in the labor supply elasticity increases wages by approximately 1.5 percents. This means that monopsony power affects negatively the wages of workers. Also, they find evidence that the elasticity of labor supply is mainly correlated with the firm effects as hypothesized in the labor economics literature. This suggests that firm market power is a key ingredient to explain heterogeneity in wage formation.

11. The sources of wage losses of displacement workers

There is scant empirical evidence on the economic mechanisms that generate the wage losses of displaced workers. The study by Raposo et al. (2019) offers a novel evaluation of the sources of wage losses incurred by workers displaced due to firm closure, bearing in mind that wages in the previous job are a function of a set of worker characteristics (for instance, gender, education, and experience) that are expected to yield, in general, the same return on the previous job and on the subsequent job, and a set of firm, job title, and match characteristics that do not necessarily yield the same return in subsequent jobs. Hence, if wages primarily reflect characteristics, then individual wages will be highly persistent and largely invariant to where individuals work, and potential losses due to displacement will be negligible. If, on the other hand, firm, job title, and match-specific heterogeneity are important, then the costs of displacement incurred by workers could be considerable.

A displacement event could lead to the loss of occupation-specific human capital due to the difficulty of finding a job that uses existing skills optimally or due to the depreciation of specific human capital during non-employment spells. Human capital has a decisive role during the early phase of the joblessness spell because larger human capital endowments are initially associated with greater job opportunities and higher opportunity costs of unemployment that necessarily erode with the progression of the unemployment spell.

Wage losses of displaced workers can be related to the firm, job title, and match quality that existed before and after displacement. The approach of Raposo et al. (2019) provides a unified framework that allows them to separately identify the components of the sources of the wage losses associated with the worker-firm pair into the contribution of worker, firm, and match quality.

In this framework it is important to distinguish a good worker in a good firm from a good worker-firm match (that is, a match with higher quality). In the event of a displacement, a loss occurs if a high-quality job match between the worker and the firm is dissolved. Furthermore, match-specific human capital accumulated over the course of the employment relationship is permanently destroyed when a job separation occurs. Its value is lost to both match participants and to the society as a whole.

Raposo et al. (2016) show that post-displacement monthly wages are, on average, 7.2 log points lower than pre-displacement wages. The results showed that sorting into job titles plays a very sizable role in explaining the losses experienced by workers displaced through firm closures, accounting for 37 percent of the total average monthly wage loss. The loss of match quality effects also plays a significant role, accounting for 32 percent of the total average monthly wage loss. The remaining 31 percent are attributed to the negative sorting of workers across firms with different pay standards.

12. The sources of the union wage gap

In a regime of near-universal collective bargaining coverage one may nevertheless discern sharp union wage gaps due to the heterogeneous influence of unions in covered settings. In the study of Addison et al. (2015) it is shown that the estimates of the union density wage gap for total monthly earnings top out at approximately 24 log points.

Employing the, by now familiar, regression approach which accounts for work, firm, job-title, match quality heterogeneity they show that the observed union wage gap would be very small if workers were randomly allocated to firms; that union workers are no more productive than their non-union counterparts; that part of the wage premium associated with union workers enjoying elevated job titles and/or benefiting from more generous promotion policies is important; and that the match quality does not play a role in the determination of the union wage gap.

The stochastic structure of estimated fixed effects can usefully be explored to understand the mechanism of wage inequality, via the use of quantile regressions. In other words, a decomposition similar to the one used for the mean can be employed to investigate the role of sorting in terms of worker skills, the generosity of firm's wage policy, and the paying conditions of different job titles.

Taking into account unobserved worker heterogeneity and sorting across firms with distinct wage policies and job titles that are differently paid, it is shown that union density contributes to less wage dispersion. But unobserved worker heterogeneity also plays an equalizing role under unionism. Unionized low productivity workers have higher wages than their non-unionized counterparts. In contrast, the allocation of unionized workers by job title makes for increasing inequality. The indication that the wage gap between unionized and non-unionized worker increases along the firm dimension is quite strong throughout the wage distribution. The contribution to wage inequality, however, is relatively mild.

13. The sources of human capital spillovers

While literature on the returns to schooling has made remarkable progress over the last 50 years the role of worker sorting across employers or jobs has been neglected. To the extent that workers with different levels of education are not randomly allocated to firms, and firms' pay standards are heterogeneous, they could be a key channel yielding returns to education. A similar argument could be built over jobs. There are remarkable wage contrasts across narrow occupations, possibly driven by differences in their degree of riskiness, the amount of specific training required, or the technology used. Provided that education can grant a "passport" to better paying jobs, part of the overall return on education would operate through a job title channel.

The study of Cardoso et al. (2018) explores the role of firm- and job-level heterogeneity shaping the returns to schooling, quantifying the impact of sorting of workers across firms and job titles on the returns to education. Analysis of the role of firm- and job-level heterogeneity structuring the returns to education begs another question: What if peers matter? Fundamentally, the quality of a firm will depend on the quality of its human resources. They progress to explicitly acknowledge that work within a firm is not undertaken in isolation, but with coworkers.

It is shown that indeed schooling grants access to better paying firms and jobs. It reveals that one fourth of the overall return on a year of education (7.9 log points) operates through the firm channel, whereas a third operates through the detailed job the worker performs. The worker component is responsible for 38 percent of the return to education.

In the second part of their analysis the authors consider the role of co-workers shaping individual wages. For this purpose, they extend the high-dimensional fixed effects framework to incorporate a measure of their peers' quality (via the average of the co-workers fixed effects). They show that peer quality has a sizeable impact driving wages. In their preferred specification, a 10 percent increase in the measure of peer quality leads to a wage increase of 2.1 log points. Very importantly, they conclude that one additional year of schooling of the co-workers would engender a half a percent increase in the wage of the worker.

14. The sources of future research

From the econometric specification point of view, the natural extension of high-dimensional fixed effect regression models is, of course, the generalisation towards varying regression coefficients, as in Gibbons et al. (2019) and Guimarães et al. (2017). In this vein, a promising empirical strategy is the transposition from linear to non-linear models, as in Félix et al. (2019). The development of theoretical models, in particular search and matching models, giving support to a valuable reduced-form interpretation of the parameters of AKM type models, as in Kline et al. (2019), is very welcome. Finally, the study of multiple and complex layers of social connections (e.g., parents, children, siblings, spouses, peers, classmates, neighbours, and friends) , as in Eliason et al. (2019)

is also seen as a recommendable way to enhance our knowledge about the influence of social networks on wages.

References

- Abowd, John M., Francis Kramarz, and David Margolis (1999). High wage workers and high wage firms. *Econometrica*, 67(2): 251-333.
- Addison, John T., Mário Centeno, and Pedro Portugal (2004). Reservation wages, search duration, and accepted wages in Europe. Bank of Portugal working paper 04-13.
- Addison, John T., Pedro Portugal, and Hugo Vilares (2017). Unions and collective bargaining in the wake of the Great Recession: evidence from Portugal. *British Journal of Industrial Relations.*, 55(3): 551-576.
- Addison, John T., Pedro Portugal, and Hugo Vilares (2015). The sources of the union wage gap: the role of worker, firm, match, and job title heterogeneity. Bank of Portugal working paper 15-12.
- Angrist, Joshua (2014). The perils of peer effects. *Labour Economics*, 30: 98-108.
- Burdett, Kenneth and Dale T. Mortensen (1998). Wage differentials, employer size, and unemployment. *International Economic Review*, 39(2): 257-273.
- Card, David, Ana Rute Cardoso, Jörg Heining, and Patrick Kline (2018). Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics*, 36(S1): S13-S70.
- Card, David, Ana Rute Cardoso, and Patrick Kline (2016). Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women. *Quarterly Journal of Economics*, 131(2): 633-686.
- Card, David, Jörg Heining, and Patrick Kline (2013). Workplace heterogeneity and the rise of West German wage inequality. *Quarterly Journal of Economics*, 128(3): 967-1015.
- Cardoso, Ana Rute, Paulo Guimarães, and Pedro Portugal (2016). What drives the gender wage gap? A look at the role of firm and job-title heterogeneity. *Oxford Economic Papers*, 68(2): 506-524.
- Cardoso, Ana Rute, Paulo Guimarães, Pedro Portugal, and Pedro Raposo (2016). The sources of gender discrimination. *Bank of Portugal Economic Studies*. 2(2):47-66.
- Cardoso, Ana Rute, Paulo Guimarães, Pedro Portugal, and Hugo Reis (2018). The returns to schooling unveiled. Bank of Portugal working paper 18-05.
- Cardoso, Ana Rute and Pedro Portugal (2005). Contractual wages and the wage cushion under different bargaining settings. *Journal of Labor Economics*, 23(4): 875-902.
- Carneiro, Anabela, Paulo Guimarães, and Pedro Portugal (2012). Real wages and the business cycle: accounting for worker, firm, and job title heterogeneity. *American Economic Journal: Macroeconomics*, 4(2): 133-152.
- Eckhout, Jan and Philipp Kircher (2011). Identifying sorting - in theory. *Review of Economic Studies*, 78(3): 872-906.
- Eliason, Marcus, Lena Hensvik, Francis Kramarz, and Oskar Nordström Skans (2019). Social Connections and the Sorting of Workers to Firms. IZA discussion paper.
- Diamond, Peter (1971). A model of price adjustment. *Journal of Economic Theory*, 3(2): 156-168.

- Félix, Sónia and Pedro Portugal (2017). Labor market imperfections and the firm's wage setting policy. Bank of Portugal working paper 17-04.
- Félix, Sónia, Paulo Guimarães, and Pedro Portugal (2019). Price stickiness: microeconomic evidence using hazard models with high-dimensional fixed effects. mimeo.
- Gelbach, Jonah B. (2016). When do covariates matter? And which ones, and how much? *Journal of Labor Economics*, 34(2): 509-543.
- Gibbons, Charles E., Juan Carlos Suarez Serrato, and Michael B. Urbancic (2019). Broken or fixed effects? *Journal of Econometric Methods*, forthcoming.
- Groshen, Erica L. (1991). The structure of the female/male wage differential: Is it who you are, what you do, or where you work? *Journal of Human Resources*, 26(3): 457-472.
- Guimarães, Paulo, Fernando Martins, and Pedro Portugal (2017). Upward nominal wage rigidity. Bank of Portugal working paper 17-02.
- Guimarães, Paulo, and Pedro Portugal (2010). A Simple feasible procedure to fit models with high-dimensional fixed Effects. *Stata Journal*, 10(4): 628-649.
- Kline, Patrick, Raffaele Saggio, Mikkel Sølvsten (2019). It ain't Where you're from, it's where you're at. Communication presented at LEED, New York.
- Manning, Alan (2003). Monopsony in motion: imperfect competition in labor markets. Princeton University Press.
- Manski, Charles (1993). Identification of Endogenous Social Effects: The Reflection Problem. *The Review of Economic Studies*, 60(3): 531-542.
- Portugal, Pedro, Pedro Raposo, and Hugo Reis (2018). The distribution of wages and wage inequality. *Bank of Portugal Economic Studies*. 4(1):1-14.
- Raposo, Pedro, Pedro Portugal, and Anabela Carneiro (2019). The sources of wage losses of displaced workers: the role of the reallocation of workers into firms, matches, and job titles. *Journal of Human Resources*, forthcoming.
- Roy, A. (1951). Some thoughts on the distribution of earnings. *Oxford Economic Papers*, 96: 201-229.
- Torres, Sónia, Pedro Portugal, John T. Addison, and Paulo Guimarães (2018). The sources of wage variation and the direction of assortative matching. *Labour Economics*, 54: 47-60.
- Woodcock, Simon (2015). Match Effects. *Research in Economics*, 69(1): 100-121.