Union membership density and wages: The role of worker, firm, and job-title heterogeneity

John T. Addison[☆a,b], Pedro Portugal[c,d], and Hugo de Almeida Vilares[e,f]

[a]*Department of Economics, Darla Moore School of Business, 1014 Greene Street, Columbia, SC 29208, United States of America*
[b]*Department of Economics and Finance, Durham University Business School, Durham University, Millhill Ln, Durham DH1 3LB, United Kingdom*
[c]*Banco de Portugal, Av. Almirante Reis 71, 1150-015 Lisboa, Portugal*
[d]*Nova School of Business and Economics, R. da Holanda n.1, 2775-405 Carcavelos, Portugal*
[e]*Centre for Economic Performance, Houghton St, London WC2A 2AE, United Kingdom*
[f]*Department of Economics and Department of Management, LSE, Houghton St, London WC2A 2AE, United Kingdom*

Abstract

We examine the association between union density and wages in Portugal where just 10 percent of all workers are union members but nine-tenths of them are covered by collective agreements. Using a unique dataset on workers, firms, and collective bargaining agreements, we examine the union density wage gap in total monthly wages and its sources — namely, worker, firm, and job-title or 'occupational' heterogeneity — using the Gelbach decomposition. The most important source of the mark-up associated with union density is the firm fixed effect, reflecting the differing wage policies of more and less unionized workplaces, which explains two-thirds of the wage gap. Next in importance is the job-title fixed effect, capturing occupational heterogeneity across industries. It makes up one-third of the gap, the inference being that the unobserved skills of workers contribute at most only trivially to the union density wage gap. In a separate analysis based on disaggregations of the total wage, it is also found that employers can in part offset the impact of the bargaining power of unions on wages through firm-specific wage arrangements in the form of the wage cushion. Finally, union density is shown to be associated with a modest reduction in wage inequality as the union density wage gap is highest among low-wage workers. This result is driven by the job-title fixed effect, low-wage workers benefiting more from being placed in higher paying 'occupations.'

JEL Classification: J31, J33, J41, J51, J52

Keywords: union density, union density wage gap, worker/firm/job-title fixed effects, Gelbach decomposition, wage inequality

[☆]Corresponding author. *E-mail addresses:* ecceaddi@moore.sc.edu (J. T. Addison), pportugal@bportugal.pt (P. Portugal), h.f.vilares@lse.ac.uk (H. de Almeida Vilares).

1. Introduction

In this paper we examine the association between union density and wages in Portugal and the role played by worker, firm, and job-title heterogeneity in that relationship. Both are interesting if challenging questions. Thus, for example, in continental Europe the architecture of wage bargaining systems typically leads to bargaining coverage rates that exceed union density rates, and often by a considerable margin. The pervasive presence of extension mechanisms that generalize union wage agreements throughout industry is responsible for this asymmetry. One might easily be led to conclude in these circumstances that unions have an indistinguishable effect on wages across firms. Vulgo: how can a union premium arise when all workers are covered?

Clearly, union membership has a different meaning in continental Europe than in the United States. In the first place, the union contract does not determine the wages of covered workers, but rather the wage floor for specific job titles. Actual wages may differ from bargained wages. Moreover, being a union member in Europe is less likely to correspond to a discrete event associated with strong union power via a (successful) union representation election at the firm, as is the case for the United States. That said, union density rates at the firm or contractual level may still convey relevant information regarding the strength of unions, union power elevating the worker share of the firm's economic rents in all theoretical models of collective bargaining. Indeed, union density may even prove a more adequate measure of the degree of dispersion in union power than in the United States. The case for measuring union power is further underscored by the existence different layers of wage negotiations in continental European nations.

With respect to what lies behind the union density wage gap, a second interesting feature of the paper is that it quantifies the importance of sorting in unionized labor markets. In particular, we examine the allocation of heterogeneous union workers across firms with heterogeneous wage policies and granular job-titles that are differently remunerated. That is, we seek answers to the following questions. Is union density systematically related to heterogeneous firm wage policies? Do we consistently observe more higher-paying job-titles (or more generous promotion policies) in more highly unionized firms? Do unionized workers possess unobserved characteristics that make them more productive? While clearly very much in the spirit of the contemporary interest in the process of allocation of heterogeneous workers to firms, this examination of the manner in which union members are assigned to firms and detailed 'occupations' has to our knowledge never before been examined in the union literature.

The contributions of the paper are as follows. Having obtained conventional linear regression estimates of the union density wage gap for total monthly wages, we then decompose the resulting value using a high-dimensional fixed effects wage regression model. Employing the Gelbach (2016) procedure, we exploit the omitted variable bias formula to decompose the union density rate regression coefficient estimate into its worker heterogeneity, firm heterogeneity, and job-title heterogeneity components. We extend the procedure by replacing the single union density coefficient estimate by a smooth nonlinear function characterization of the union effect. We then seek to peer more closely inside the mechanisms of wage setting, distinguishing between the bargained wage determined at sectoral level and the wage cushion, where the latter is defined as the difference between the bargained wage and total wages. Lastly, we estimate a three-way high-dimensional quantile regression model, using the Machado and Santos Silva (2019) method of moments estimator, to address the potential impact of union power on wage inequality and, no less important, its component sources.

Our main findings with respect to total wages are fivefold. First, the union density wage gap is around 16 log points on average. Second, the influence of union density is not linear. Some critical mass of union density — around 30 percent — is required to have any material influence on wages, while beyond some level (approximating 70 percent) further increases in union density detract from the peak premium. Third, the union density wage gap would change profoundly were workers randomly assigned to firms and job titles; that is, firm heterogeneity accounts for two-thirds of the gap. Job-title heterogeneity, whereby unionized workers are assigned to higher paying job titles in the compensation tables of collective agreements, accounts for the remaining one third, so that compensation for their unobserved ability plays a negligible role in determining the premium enjoyed by union workers. Fourth, the impact of union density on wage dispersion though of modest dimension is inequality reducing. Thus, workers in the 10th percentile record a higher wage gap than their counterparts in the 90th percentile (17.2 and 15.0 log points, respectively). After taking into account sorting and worker heterogeneity, it emerges from our extension of the Gelbach procedure that the wage compressing effect of union density mainly stems from low-wage workers benefiting more from being placed (or promoted) into higher paying job titles than is the case for high-wage workers. Finally, returning to the overall picture, we conclude with all necessary caution that the union density wage gap is likely to be between 6 and 16 log points. We base this interpretation on the notion that job-title fixed effects largely capture the effect of trade unions. A postscript from our work on the wage cushion is that in sectors where unions are weaker, relatively low bargained wages offer scope for compensatory adjustment at firm level.

2. The bargaining framework

Collective bargaining in Portugal occurs mainly at the sectoral or industry level, where wage ordinances only set wage floors. Further, although in a strict sense these union negotiated wage agreements only formally bind the members of the employer's association and the affiliated workers of the signatory trade union(s), there is in practice a widespread and well anticipated practice of legal extensions that dramatically broaden the incidence of the agreements. The upshot of this near automatic procedure is that every firm within the sector in which the agreement is extended, and every worker within those firms, are no less covered than the original signatories themselves.[1]

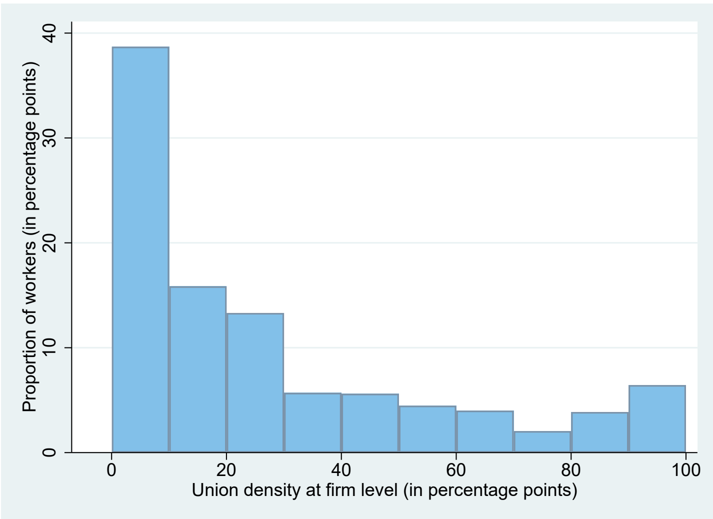

Fig. 1. The distribution of union density conditional on positive union density. Source: Relatório Único, 2010-2013.

[1]Individual companies and one or more unions may sign firm-level agreements, known as *Acordos de Empresa* (or AEs), which represent a better-paying alternative to the nearest-neighbor sectoral agreement. Their prevalence is limited to less than 5 percent of the private sector workforce, and their existence is mostly related to idiosyncrasies in the firms' shareholder structures. A similarly modest proportion of workers is covered by agreements established between a group of firms and one or more unions under *Acordos Colectivos de Trabalho* (or ACTs).

This decoupling of coverage and membership, both within and between firms, translates into levels of collective bargaining coverage in the order of 90 percent of workers at a time when union density approximates 10 percent. That is, around 80 percent of the labor force have benefited from collective agreements without being members of a trade union. It further explains the finding that around 62 percent of workers are in firms with zero union density. And among those in firms with positive union density, the distribution is highly skewed toward low values. As shown in the histogram of Fig. 1, close to two-fifths of these observations record less than 10 percent density.[2]

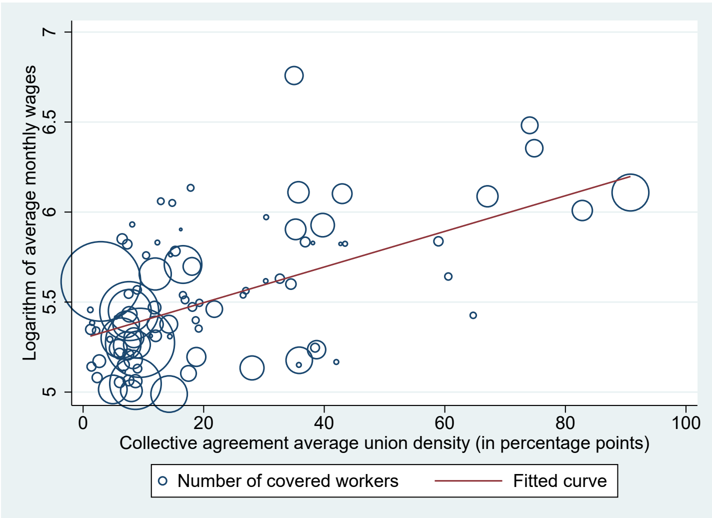

Fig. 2. Collective agreement size, average total wage at collective agreement level, and average union density at firm level. Source: Relatório Único, 2010-2013.

As a result, while individual membership of unions and employers' associations becomes inconsequential for collective bargaining agreement coverage purposes, it does not follow that membership is inconsequential. It is unlikely that a union representing 2 percent of the workers in an industry will have the same success as a union representing 40 percent of the workers, even if both contracts record 100 percent coverage. In other words, heterogeneity in union power (as proxied by union density) should after all be reflected in wage differentials. This is indeed the indication of Fig. 2 showing that union

[2]The average union density rates changed from 10.3 percent in 2010, to 9.7 percent in 2011 and 2012, and to 9.9 percent in 2013. During the 2010-2013 period, the between component of the union density variance is 0.0416 and for the within component is 0.0057. More generally, for analysis of the effect of union decline on wage inequality see DiNardo et al. (1996), Card et al. (2004), and Farber et al. (2021).

density rates comove with wages. In summary, union membership in continental Europe is best seen as a continuous process measuring union power rather than a discrete event always identifying high union densities.

Table 1
The hospitality sector collective agreement, 2017

Rank	Job Description	Minimum Monthly Base Wage (in Euros)
13	General Manager	1,515
12	Board Assistant; Commercial Manager; Service Manager; Human Resources Manager; Technical Manager	1,240
11	Head of Department; Head of Division; Head of Services; Nutrition Technician 1st Class	1,018
10	Head of Section (office); Head of Sales; Inspector; Board Secretariat Officer; Nutrition Technician 2nd Class	898
9	Administrative; Head of Cafeteria; Head of Purchases; Head of Kitchen; Head of Pastries; Head of Storage; Head of Dinning Room; Inspector of Sales	808
8	Cashier; Head of Preparation Room; Controller; Cook of 1st Class; Sub-Head of Dinning Room; Administrative Assistant; Pastry Cook; Sales Technician	771
7	Driver of Heavy Vehicles; Storage Keeper; Polyvalent Worker	716
6	Driver of Non-heavy Vehicles; Administrative Assistant 2nd Class; Pastry Cook 2nd Class; Sub-Head of Dinning Room 2nd Class; Sales Representative	700
5	Cook 2nd Class; Controller of Balcony; Controller of Bar; Controller of Storage and Packing; Admin. Assistant 3rd Class	629
4	Head of Copa; Cook of 3rd Class; Packing Worker; Storage Worker	582.50
3	Controller Cashier; Storage Worker; Bar Worker; Balcony Worker 1st Class; Distribution Handler;	570
2	Balcony Worker 2nd Class; Admin. Intern; Hospitality Assistant	562
1	Driver Assistant; Distribution Assistant; Barman Intern (1 year); Cook Intern (1 year); Pastry Intern (1 year); Cleaning Worker; Dining Room Employer	557

Notes: Table extracted from the Sectoral Agreement of April 22, 2017, between *AHRESP* (the Association (of Employers) in Hotels, Restaurants and Similar) and *SITESE* (the Union of Workers and Technicians of Services, Commerce and Hospitality). Source: Boletim do Trabalho e Emprego, 2017.

In a more granular setting, the strength of unions can be expressed in negotiations conducted at the industry level mainly through the determination of industry wide wage floors, or bargained wages, for specified job titles ("categorias profissionais"); the definition of job titles themselves over which wage floors are negotiated; and the promotion rules that automatically place workers into higher job titles. The partitioning of the workforce into these occupational categories, whose specificity is sufficient to justify bargaining over, takes into consideration the complexity of the tasks involved, the hierarchical standing of the worker, and the demands of the associated working conditions. In a typical year, the firmament of 300 sectoral or industry level agreements, each containing fully specified tables of job titles, constitutes the core outcome of collective bargaining.[3] Some 30,000 different job titles/bargained wages are prescribed market wide. Table 1 illustrates by means of a representative agreement with its associated complement of job titles and bargained wages.[4]

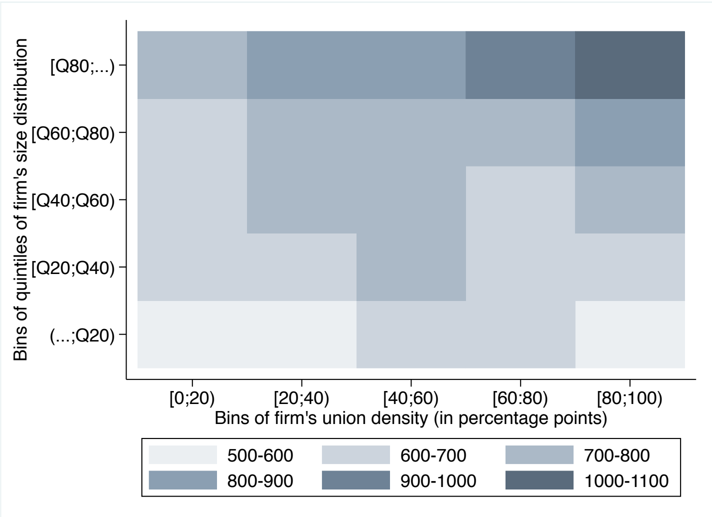

Fig. 3. Heatmap of union density, firm size, and average real total wages (in Euros) at firm level. Source: Relatório Único, 2010-2013.

While exhaustive, these sectoral wage tables neither fully determine wages nor completely define their application to individual workers. At firm level, the power of trade unions may be manifested in higher compliance with collective wage agreements; in negotiated wages exceeding those agreed at the industry level; in more favorable promotion rules; and by new or more generous productivity or tenure related premia. Consequently, firms frequently pay more than bargained wages, particularly larger and more unionized firms. In sum, wage negotiations at firm level are expected to generate wage gains above and beyond those obtained at industry level. Indeed, this is the suggestion of Fig. 3

[3]For a detailed review of recent developments in collective bargaining and extension arrangements in Portugal, see Hijzen and Martins (2016), Addison et al. (2017), and Card and Cardoso (2021).

[4]Table A.1 provides some details on the largest 15 collective agreements, including the number of job titles specified per agreement and the average union density of the constituent firms.

which shows that firm level union density rates impact differently in firms of different sizes. More concretely, union wage effects are higher among larger firms.

3. The datasets

The data sources used in this exercise are the *Quadros de Pessoal* (Personnel Tables), from its inception in 1986 until 2009, and the successor *Relatório Único* (Single Report) from 2010-2013.[5] For our purposes, each longitudinal matched employer-employee-contract-job title database is identical other than in one main respect: the follow-up survey contains data on the union density of the firm that for the first time permit accurate estimates of union density to be obtained. We will therefore not distinguish between the two other than to address the major, albeit temporary, innovation in the successor survey, and to note that while the union density wage gap can only be computed using the *Relatório Único*, both surveys are instrumental in computing the conditional decomposition of estimates of the union density wage gap in to its component firm, worker, and job-title fixed effects.

The surveys are mandatory in nature and are administered by the Ministry of Employment and Social Security on an annual basis for all establishments with at least one wage earner. All workers employed by the firm in the reference month (March of each year until 1993, October thereafter) are reported, although civil servants and workers in domestic service are not covered while the coverage of agriculture is necessarily uneven because of the importance of informal work/low share of wage earners in this sector. In short, the entire population of private-sector firms in manufacturing and services with wage earners is covered. Further, by virtue of their mandatory nature, the high response rate in the surveys ensures that problems commonly associated with panel data such as panel attrition are much attenuated. Confidence in the coverage and reliability of the surveys is further underscored by the requirement that the data be made publicly available at the place of work.

The databases report the location, industry, employment, sales, ownership, and legal basis of the establishment/firm. Worker information includes gender, age, skill, broad occupation, schooling completed, starting date with the firm, earnings, and working

[5] For the years 1990 and 2001 the *Quadros de Pessoal* survey is not available. Also, although the union question is ongoing, data on union density in the *Relatório Único* have not been made available to researchers since 2013.

hours. In addition, the surveys also record the collective bargaining arrangement and the specific job title held by the worker defined in the collective agreement. Recall that this job classification variable goes beyond a fine definition of occupations to encompass the complexity of the task performed, the skill level of the job, the required labor market experience, and the hierarchical standing of the worker. The wage variable records the worker's gross monthly earnings (the actual or total wage as of the survey month), which sum is split into the following four components: the base wage (i.e. the gross pay for normal hours of work), overtime pay, and regularly and irregularly paid supplements. Normal monthly hours worked and overtime hours are also reported.

The Portuguese surveys are administered exhaustively to all wage earners in the entrepreneurial sector.[6] That is, the surveys include all workers whose contracts are covered by collective agreements as well as all workers that are not covered. Non-covered workers are identified as such whenever there is no collective agreement assigned to them, one consequence of which is that information on a worker's job-title will be missing. The main reason for non-coverage is simply the absence of a wage agreement in a particular industry, occupation, or firm. Also, in some cases, a failure to renew a collective agreement beyond its term may lead temporarily to non-coverage status.

In the current study, we decided to exclude non-covered workers. As a result, over the 2010-2013 period, 743,884 (or 9.7 percent of) observations, corresponding to 320,464 (or 11.4 percent of) workers, and 25,508 (or 8.3 percent of) firms were excluded. The subsample of non-covered workers is both heterogeneous and unrepresentative, and it is dominated by those in real estate (25.2 percent) and the non-profit health and social services industry (21.7 percent).[7]

The following restrictions were placed on the data. First, the analysis was confined to full-time employed workers in receipt of what was contractually defined for the survey reporting month. Second, workers from the agriculture, fisheries, and energy products/extraction sectors were excluded. Third, workers aged less than 18 years and more than 65 years were excised, as also were those whose monthly wages were less than 80

[6]Public administration, where union density rates are conspicuously higher than in the private sector, is not covered in the *Quadros de Pessoal* and *Relatório Único*. There are no official statistics on union density among civil servants but is widely believed it may exceed 50 percent in public health and public education sectors.

[7]In Table A.2 it is shown that the inclusion of non-covered workers does not impact significantly the regression coefficient estimate for union density in the base model. Estimation of the full model is, of course, not possible for non-covered workers.

percent of the mandatory minimum wage, corresponding to the lowest admissible wage for apprentices. Finally, observations not belonging to the largest connected group were dropped, amounting to some 1 percent of the total number of observations.

This brings us to the distinguishing feature of the successor survey — the *Relatório Único* — namely that it allows us to construct a measure of union density at firm level. Specifically, the survey asks of the manager respondent: *Indicate the number of workers for whom you have knowledge of their membership in a union (because they are union officials, because you deduct membership dues from their salary, or because the worker informed you about his/her membership so as to determine which particular collective regulation is applicable to their case).* The sum total of such workers — whose personal union status is unknown, thereby precluding use of an individual union membership variable — divided by the number of workers employed by the firm provides our measure of union density.

Overall, the joint dataset includes 36,616,379 observations of worker-year pairs, of which 6,218,777 are from the *Relatório Único*. The joint dataset has a basis in the records of 6,042,315 workers matched by identifying social security number, 652,487 firms matched by identifying number, and 132,908 job-titles matched by the code of the collective agreement occupational category. Table A.3 provides the descriptive statistics for observations from both the *Quadros de Pessoal* and the *Relatório Único*.

4. Modeling

We next describe the procedures used to estimate the union density wage gap and to account for the component contributions of firm compensation policies, worker ability, and detailed occupational premiums (via the estimation of firm, worker, and job-title fixed effects, respectively).

4.1. Estimation of the union density wage gap

We begin with a standard Mincerian wage equation, augmented to include union density, as follows:

$$w_{it} = \boldsymbol{x'_{it}}\boldsymbol{\beta_0} + \delta_{0_t} + \gamma_0 U_{F(i,t)} + \epsilon_{0_{it}}, \tag{1}$$

where w_{it} is the natural logarithm of worker i monthly compensation at year t, $\boldsymbol{x'_{it}}$ is a vector of observed characteristics of the worker and his/her employer, $\boldsymbol{\beta_0}$ is a vector of coefficients for the observed characteristics of workers and firms, $U_{F(i,t)}$ is the level of union density of employer F in year t, γ_0 is the coefficient associated with the level of union density, δ_{0_t} are calendar year fixed effects included to capture the business cycle, and $\epsilon_{0_{it}}$ is an error term, assumed to be uncorrelated with the covariates. The explanatory variables (or observed characteristics of workers and firms) are age, age squared, seniority, seniority squared, and dummies for gender, education, firm size, and industry.[8]

To allow for a non-proportional impact of union density on wages we will consider both parametric and non-parametric approaches. The parametric version employs a third-degree polynomial in union density:

$$w_{it} = \boldsymbol{x}'_{it}\boldsymbol{\beta_0} + \delta_{0_t} + \gamma_{0_1} U_{F(i,t)} + \gamma_{0_2} U^2_{F(i,t)} + \gamma_{0_3} U^3_{F(i,t)} + \epsilon_{0_{it}}. \tag{2}$$

In the non-parametric version, the impact of union density is captured by the presence of fixed effects $\psi_{0u_{F(i,t)}}$ (one for each different level of union density U):

$$w_{it} = \boldsymbol{x}'_{it}\boldsymbol{\beta_0} + \delta_{0_t} + \psi_{0u_{F(i,t)}} + \epsilon_{0_{it}}. \tag{3}$$

Information contained in the union fixed effects while necessarily complete is rather noisy and "staccato." Thus, in a second step, we will estimate a kernel regression linking the estimates of the union density fixed effects and actual union density at firm level, as follows:

$$\widehat{\psi_u} = K(U) + \upsilon_u, \tag{4}$$

where $\widehat{\psi_u}$ is the union density fixed effect estimate obtained from the first step, U is the prevailing union density of the firm, υ_u is the disturbance term, and K is a standard Epanechnikov kernel function (Silverman, 1986).

The estimation of local weighted union density wage gaps, as well as the third-degree polynomial specification, result in smoothed estimates of a union density wage gap curve in actual earnings. To facilitate interpretation of the results, a convenient normalization in the nonparametric case requires that the fixed effect in the absence of workplace unionism be set equal to zero. No further restrictions are implied by this assumption as the union density wage gap represents the relative difference in wages

[8]The subscript **0** denotes the base regression model specification.

for workers at firms with different levels of union density, controlling for the observed characteristics of workers and firms.

4.2. Estimation of the sources of the union density wage gap

Given the estimate of the union density wage gap it is useful to decompose this outcome measure into its constituent mechanisms; that is, to identify the contributions of worker, firm, and job-title time-invariant heterogeneity.

For the full-specification model, we include in equation (3) the sources of time-invariant heterogeneity, namely the worker fixed effect (α_{1_i}), the firm fixed effect ($\lambda_{1_{F(i,t)}}$), and the job-title fixed effect ($\theta_{1_{J(i,t)}}$), exploiting the methodology first introduced in Carneiro et al. (2012). The model thus becomes:[9]

$$w_{it} = \boldsymbol{x'_{it}\beta_1} + \delta_{1_t} + \gamma_1 U_{F(i,t)} + \alpha_{1i} + \lambda_{1_{F(i,t)}} + \theta_{1_{J(i,t)}} + \epsilon_{1_{it}}. \tag{5}$$

In general the identification of the worker, firm, and job title fixed effects is assured by the restriction that the sample identifies the largest connected set. A connected set is defined when at least one element of a worker, firm, and job title combination links the rest of the group (Abowd et al., 1999). The largest connected group represents more than 99 percent of the sample.

At this stage, we calculate the independent contribution of each fixed effect to the union density wage gap. For this purpose we adapt the methodology developed in Gelbach (2016), which appeals to the omitted variables bias formula to compute a detailed decomposition. Departing from a baseline specification to which covariates are added, Gelbach's procedure allows us to compute the contribution of each new covariate to the change in the estimate of the coefficient of the variable under scrutiny. In our case, it allows us to unambiguously disentangle the contribution of each excluded variable (each fixed effect) to the variation of the coefficient estimate(s) of the union density variable(s).

To better understand our decomposition exercise it is useful to present the benchmark wage regression equation, emphasizing the union effects, in the following matrix formulation:

$$\mathbf{W} = \mathbf{X}\boldsymbol{\beta_0} + \mathbf{U}\boldsymbol{\gamma_0} + \boldsymbol{\epsilon_0}, \tag{6}$$

[9]The subscript **1** denotes the full model specification.

where $\mathbf{W}$ stands for vector of wages, $\mathbf{X}$ denotes the matrix of control variables, including the year dummies, $\boldsymbol{\beta_0}$ is a vector of regression coefficients, $\mathbf{U}$ collects the union density variable(s), $\boldsymbol{\gamma_0}$ represents the union density wage gap, and $\boldsymbol{\epsilon}$ is the vector containing the error terms.

Making use of the Frisch-Waugh-Lovell theorem, we can express the OLS estimate of $\boldsymbol{\gamma_0}$ by running a regression of $\boldsymbol{W}$ on $\mathbf{U}$ after partialing out the effect of $\mathbf{X}$ on both variables. That is,

$$\widehat{\boldsymbol{\gamma_0}} = (\mathbf{U}'\mathbf{M_X}\mathbf{U})^{-1}\mathbf{U}'\mathbf{M_X}\mathbf{W}, \tag{7}$$

where, $\mathbf{M_X} = I - X(X'X)^{-1}X'$ is the residual-maker, or "annihilator" matrix.

More compactly, we can write:

$$\widehat{\boldsymbol{\gamma_0}} = \mathbf{A_X}\mathbf{W}, \tag{8}$$

and introduce the definition of the matrix $\mathbf{A_X} = (\mathbf{U}'\mathbf{M_X}\mathbf{U})^{-\mathbf{1}}\mathbf{U}'\mathbf{M_X}$, which will be instrumental in the application of the omitted variable bias formula. In general, if we pre-multiply any variable by $\mathbf{A_X}$ we will always obtain the corresponding regression coefficient estimate of the union density variable, after controlling for the variables included in $\mathbf{X}$.

We now define the the full regression model, where we incorporate the worker effects (identified via the matrix $\mathbf{D}$), the firm effects (identified via $\mathbf{F}$), and the job-title effects (identified via $\mathbf{J}$). The estimated full regression can be now expressed as:

$$\boldsymbol{W} = \mathbf{X}\widehat{\boldsymbol{\beta_1}} + \mathbf{U}\widehat{\boldsymbol{\gamma_1}} + \mathbf{D}\widehat{\boldsymbol{\alpha_1}} + \mathbf{F}\widehat{\boldsymbol{\lambda_1}} + \mathbf{J}\widehat{\boldsymbol{\theta_1}} + \widehat{\boldsymbol{\epsilon_1}}, \tag{9}$$

where $\widehat{\boldsymbol{\alpha_1}}$, $\widehat{\boldsymbol{\lambda_1}}$, and $\widehat{\boldsymbol{\theta_1}}$ denote the worker, firm, and job-title fixed effects, respectively.

At this stage, we build on the approach suggested by Gelbach (2016), which makes use of the OLS omitted variable bias formula, to decompose the union density wage gap in terms of individual self-selection in unionized firms (the worker component) and sorting across firms with different wage policies and differently remunerated job titles. This can be achieved by multiplying both terms of the full regression by $\mathbf{A_X}$, leading to:

$$\widehat{\boldsymbol{\gamma_0}} - \widehat{\boldsymbol{\gamma_1}} = \mathbf{A_X}\mathbf{D}\widehat{\boldsymbol{\alpha_1}} + \mathbf{A_X}\mathbf{F}\widehat{\boldsymbol{\lambda_1}} + \mathbf{A_X}\mathbf{J}\widehat{\boldsymbol{\theta_1}} = \widehat{\boldsymbol{\tau_{\alpha_1}}} + \widehat{\boldsymbol{\tau_{\lambda_1}}} + \widehat{\boldsymbol{\tau_{\theta_1}}} \tag{10}$$

as, by construction, $\mathbf{A_X}\widehat{\boldsymbol{\epsilon_1}} = \mathbf{0}$.

Equation (10) yields an exact, unambiguous and conditional decomposition of the union density wage gap. The interpretation of this equation is that we can split the wage gap into three components: a worker component ($\widehat{\tau_{\alpha_1}}$), a firm component($\widehat{\tau_{\lambda_1}}$), and a job-title component ($\widehat{\tau_{\theta_1}}$). In practice, all we need to do is to run a regression for each type of fixed effect on all regressors of the benchmark regression (**X** and **U**) and extract the union regression coefficient estimates.

Mutatis mutandis, we can apply the same principle of the Gelbach decomposition to the union density wage gap curve given in equation (2) and to the union fixed effect specification in equation (3). In the latter, the difference between the union density fixed effects of the full and base models can be decomposed into three fixed effects:

$$\widehat{\boldsymbol{\psi}_{\mathbf{0}_u}} - \widehat{\boldsymbol{\psi}_{\mathbf{1}_u}} = \mathbf{A}_{\mathbf{X}_\mathbf{u}}\mathbf{D}\widehat{\boldsymbol{\alpha}_{\mathbf{1}_u}} + \mathbf{A}_{\mathbf{X}_\mathbf{u}}\mathbf{F}\widehat{\boldsymbol{\lambda}_{\mathbf{1}_u}} + \mathbf{A}_{\mathbf{X}_\mathbf{u}}\mathbf{J}\widehat{\boldsymbol{\theta}_{\mathbf{1}_u}} = \widehat{\tau_{\alpha_u}} + \widehat{\tau_{\lambda_u}} + \widehat{\tau_{\theta_u}}, \tag{11}$$

where the subscript u is used to emphasize that we are decomposing the changes in the union density fixed effects. In practice, and as before, the decomposition is achieved by estimating three auxiliary regressions in which the worker, firm, and job-title fixed effects become the dependent variables and the regressors match those of equation (3). Then, by smoothing these estimates via a kernel function, we can provide a graphical representation of the components of the union density wage gap.

Before turning to our empirical results, however, we should resist the notion that the union density fixed effect is simply to be equated with a firm fixed effect. Even if union density were not to change over time, the union density fixed effect is to be seen as subsumed in the firm fixed effect in the same way that the gender fixed effect is subsumed in the worker fixed effect. Contrary to intuition, this fact does not preempt the decomposition of the union density effect along its worker, firm, and job-title dimensions for the same reason that the gender effect can be disentangled along the worker, firm, and job-title dimensions (Cardoso et al., 2016).[10] A clear indication that the dominant role of the firm is not mechanically implied by the fact that the union density variable is computed at the firm level will be given by the decomposition exercise below.

[10] Cardoso et al. (2016) extend the Gelbach procedure and prove that it is applicable even in settings where the omitted variables are fixed effects and the coefficient under scrutiny refers to a variable (viz. gender) that is subsumed in one of the fixed effects.

5. Main findings

5.1. The union density wage gap curve for monthly wages

In Portugal evidence of sizable wage differentials associated with a firm's degree of unionization is unmistakable, even though nearly every worker benefits from union bargaining. For example, the heuristic distributions of the logarithm of total hourly wages shown in Fig. 4 display meaningful differences in both shape and mean when unionized and non-unionized workplaces are considered.[11]

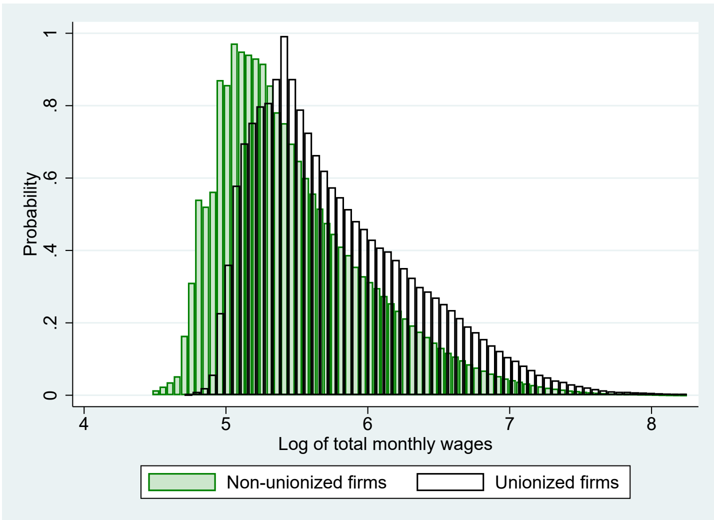

Fig. 4. Distribution of total monthly wages by union status. Source: Relatório Único, 2010-2013.

These stylized facts are confirmed by our benchmark results presented in Table 2 which chart the impact of union density on wages after having controlled for the full set of observed worker and firm characteristics, as described in section 4. For the linear specification given in the first column of the table, the estimated union density wage gap is 17.6 percent $[(e^{0.1619} - 1) \times 100]$; a sizable union wage differential that is either on a par with or exceeds U.S. estimates.[12] This wage gap is to be interpreted in the following way: it represents the wage difference between two observationally identical workers, one of whom is employed in a fully unionized firm and the other in an otherwise identical non-unionized firm.

[11]The values of zero and greater than zero were chosen because of the large majority of covered firms without union members.

[12]See the early studies of Blanchflower and Bryson (2003), and Hirsch (2004); and, especially, the more recent plant-level studies of Frandsen (2012) and Lee and Mas (2012).

Table 2
OLS estimation of the union density wage gap for total monthly wages

Variable	Linear Specification	Polynomial Specification
Union density	0.1619	−0.3844
	(0.0244)	(0.1474)
Union density squared	-	1.9030
		(0.5533)
Union density cubed	-	−1.4390
		(0.4290)
Worker's age	0.0281	0.0282
	(0.0006)	(0.0006)
Worker's age squared	−0.0002	−0.0002
	(0.0000)	(0.0000)
Tenure of the worker	0.0156	0.0156
	(0.0005)	(0.0005)
Tenure of the worker squared	−0.0002	−0.0002
	(0.0000)	(0.0000)
Female	−0.2109	−0.2108
	(0.0035)	(0.0034)
Primary school	0.0548	0.0550
	(0.0058)	(0.0056)
Basic school	0.1586	0.1581
	(0.0062)	(0.0060)
Elementary school	0.2670	0.2662
	(0.0068)	(0.0065)
Secondary school	0.4269	0.4249
	(0.0091)	(0.0086)
Post-secondary school	0.5569	0.5550
	(0.0140)	(0.0139)
University attendance	0.8143	0.8128
	(0.0117)	(0.0111)
College degree	0.9025	0.8995
	(0.0104)	(0.0099)
Firms with 50 to 99 employees	0.1316	0.1372
	(0.0041)	(0.0042)
Firms with 100 to 499 employees	0.1708	0.1817
	(0.0066)	(0.0069)
Firms with 500 to 999 employees	0.1929	0.2043
	(0.0190)	(0.0190)
Firms with 1000 to 4999 employees	0.1561	0.1700
	(0.0203)	(0.0202)
Firms with more than 5000 employees	0.1171	0.1400
	(0.0478)	(0.0475)
R^2	0.4847	0.4859

Notes: The dependent variable is total monthly wages (in logs). The controls also include 25 sector of activity dummies, and 3 year dummies. The number of observations is 6,218,777. Robust firm clustered standard errors are in parentheses. All coefficients are statistically significant at the 0.01 confidence level. Source: Relatório Único, 2010-2013.

Panel A. Polynomial union density wage gap curve

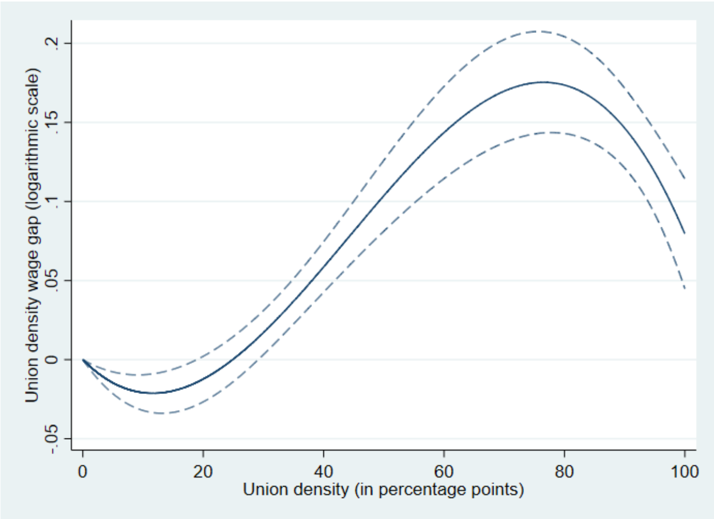

Panel B. Kernel smoothed union density wage gap curve

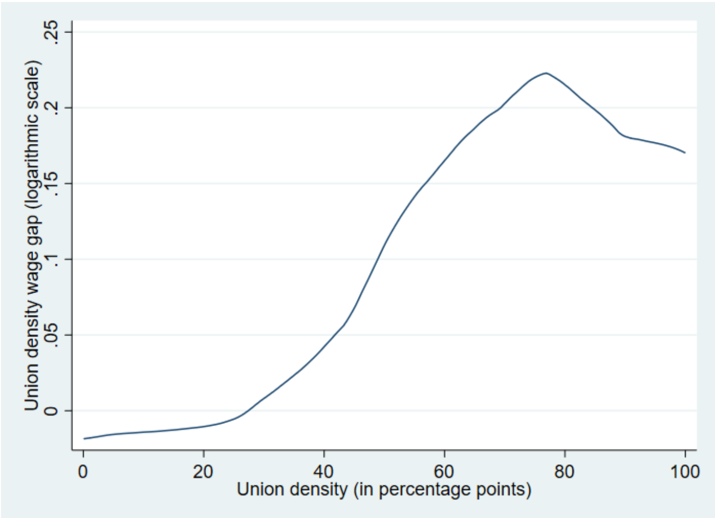

Fig. 5. The union density wage gap curve for total monthly wages. Notes: The base model includes as regressors a quadratic term in age, a quadratic term in tenure, schooling dummies (7), a gender dummy, firm size dummies (5), year dummies (3), and sector dummies (25). Robust clustered firm-year standard errors were used. In Panel A, the 95 percent confidence interval is indicated. Source: Relatório Único, 2010-2013.

The preceding methodology implies that the value of the union density wage gap for each point in the continuum of union density is determined by and conforms to a linear relationship. However, an important issue is whether the marginal change in the union density wage gap is in fact the same when a newly unionized worker joins a union-free workforce as opposed to a situation in which, say, a plurality of workers is already organized. In seeking to estimate a union density wage gap without assuming constant marginal effects throughout, we shall follow the two procedures described earlier to estimate the union density wage gap curve.

These estimates are shown in Fig. 5. Clearly, the linear approach is misleading; in particular, unions need some critical mass (of unionized workers) in order to materially influence wages. Panel A indicates that union density attains statistical significance at around 30 percent, with a maximum wage gap of 17 log points being achieved once union density reaches approximately 70 percent (see also the first three rows of the second column of Table 2). Panel B of the figure shows that although the polynomial is a sensible parsimonious approximation to the wage gap curve it understates the peak premium (now in excess of 24 log points) and overstates the decline in the premium thereafter.

The explanation for the importance of union density to wage setting has to do with the intensive margin of representation. The mere realization that a bargaining instrument covers a given worker does not seem to shed sufficient light on the properties of an agreement, namely the specific environment in which it was agreed.

The importance of considering the intensive margin is also implicit in the shape of our union density wage gap curves. A plausible explanation for the configuration of the fitted curves relies on the idea that the bargaining power of a union is a function of its ability to credibly threaten the employer through a withdrawal of labor (e.g. Farber, 1986). It is reasonable to assume that unions need some minimum complement of unionized workers to effectively impose costs on the employer in the event of a failure to agree. With a preponderance of the workforce organized, the capability to impose a total shutdown is implied, such that further increases in union density are not to be equated with higher union wage premia.

5.2. The sources of the union density wage gap for total monthly wages

The union density wage gap coefficient estimates for total earnings provide an average differential between the wages of two observationally identical workers in two observationally identical firms with distinct levels of unionization. What are the potential sources of this sizable union density wage gap? As leading contenders, we next consider the contributions of heterogeneity in the compensation policies of firms, the rules governing how the workforce is assigned to the compensation tables of the collective agreement, and the allocation of workers of different unobserved ability.

In principle, the conditional influence of unions on earnings compensation can arise from sources other than these. However, to anticipate one of our key findings, we report that after accounting for firm, job-title, and worker fixed effects the portion of the union density wage gap remaining to be explained is vestigial. This is the case for both the linear approach and the fitted union density wage gap curve. In decomposing the gap, therefore, our focus will be upon the contributions of each of these three sources of unobserved heterogeneity. In what follows, the major difference between the two (decomposition) approaches resides in the flexibility of the estimates, namely the improved estimation offered by the union density wage gap curve over the restrictive linear approach.

Job title refers to the worker's assigned role at the firm, as explicitly defined in the collective sectoral agreement governing the employment relationship. This defined "occupation" most importantly determines a floor for the base wage that a worker is legally entitled to receive. Note that the base wage set at sectoral level (which we call the *bargained wage*) does not necessarily equal the actual base wage paid by the firm. Indeed, a majority of firms pay more than the bargained wage. The difference between the base wage and the base wage floor or bargained wage is essentially determined at the discretion of the firm.

Thus, the job-title fixed effect summarizes the influence of the compensation floor defined for each worker. Each job title is collective agreement specific, such that two workers performing the same task and having the same responsibilities covered by different collective bargaining arrangements will be assigned different job titles. Overall, in a given year there are around 30,000 collective agreement/occupational category combinations to which workers are assigned. The inclusion of job-title fixed effects may be viewed as building upon a first generation Mincerian wage equation that recognizes only a broad definition of job descriptions.

Panel A of Fig. 6 shows the empirical distribution of the job-title fixed effects, contrasting workers in 'union' firms with their counterparts in 'nonunion' firms. For the former, the distribution of job-title fixed effects is visibly displaced to the right. The implication is that that better-paid job titles tend to be more heavily populated by workers of unionized firms (after taking into account firm and worker heterogeneity).

For its part, the firm fixed effect captures the (constant) wage policy of the firm,

Panel A. Job-title fixed effect empirical distribution

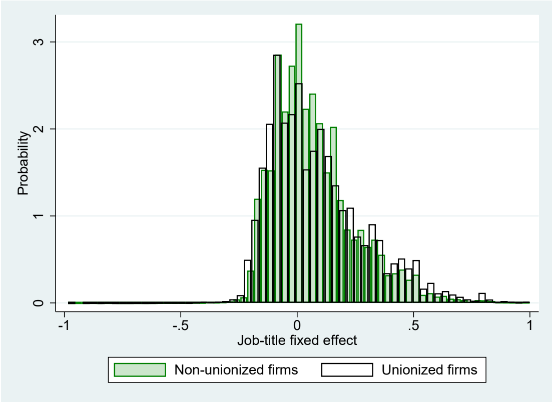

Panel B. Firm fixed effect empirical distribution

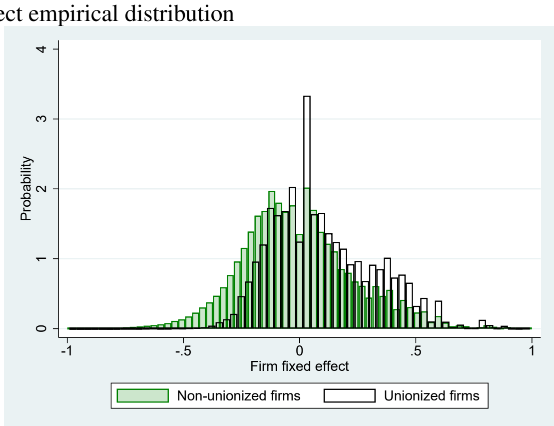

Panel C. Worker fixed effect empirical distribution

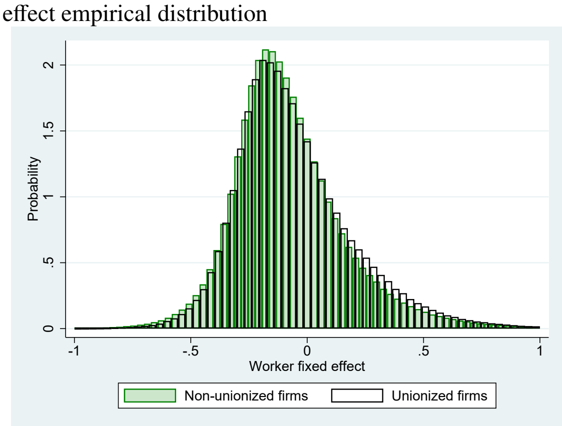

Fig. 6. Distribution of worker, firm, and job-title fixed effects by union status. Notes: In addition to the fixed effects, the model includes as regressors a quadratic term in age, a quadratic term in tenure, schooling dummies (7), a gender dummy, firm size dummies (5), and sector dummies (25). Sources: Quadros de Pessoal 1986-2009; Relatório Único, 2010-2013.

including the relative standing of the firm's wage tables, after having controlled for the placement of workers into the distinct job categories presented in such tables, which procedure fully captures the previously discussed job-title fixed effect. Firms with generous compensation policies will exhibit positive firm fixed effects, while firms with compensation policies close to the bargained wage will generate negative fixed effects. In Panel B of Fig. 6 we contrast the distribution of the firm fixed effects for workers in union and nonunion firms.[13] Clearly, unionized workers disproportionately populate high-paying firms.

Finally, the empirical distribution of the worker fixed-effects is presented in Panel C of Fig. 6. These effects capture the influence of the constant characteristics of individuals on their wages. They are essentially a proxy for the portable human capital (or productivity) of the worker. The pattern revealed is one in which more unionized firms seemingly employ relatively more skilled individuals. This outcome can be the result of observed characteristics (such as schooling or gender) or unobserved factors (ability), and we shall subsequently address the specific role of the latter.

Table 3 presents the results of the Gelbach decomposition for the linear specification.[14] The coefficient estimate contained in the first column of the table simply recalls the estimated union density wage gap (of 16.2 log points) obtained from equation (1) and reported earlier in Table 2. The estimated union density wage gap, after the inclusion of the three high dimensional fixed effects (equation (5)), is no longer statistically significant and is given in the second column of the table (-0.3 log points). To attenuate the bias arising from the incidental parameter problem in the estimation of high-dimensional fixed effect regression models, we employ the full sample, 1986-2013. To guarantee that the results are comparable with those provided in Table 2, we interact all the explanatory variables with a time dummy identifying the period before and after 2010.[15]

The third column provides the contribution of each fixed effect to the change in the estimated union density wage gap (equation (10)).[16] The differences in firms' compensation policies explain a large fraction of the union density wage gap. After accounting

[13] Observe, however, that in this comparison the influence of variables such as industry or firm size is still subsumed in the firm fixed effect. The subsequent Gelbach decomposition will enable us to filter out the impact of the firm fixed effect on the wage gap from the variable included in the benchmark specification.

[14] In Table A.4 we present the decomposition results for alternative sets of observable controls.

[15] Notice that the use of these interaction terms, in the case of the base model, guarantees that the regression coefficient estimates for the 2010-2013 period are exactly the same as those given in Table 2.

[16] In practice, as we have seen, the application of the Gelbach decomposition in the current framework amounts to running three auxiliary OLS regressions identical to equation (1), but where the dependent

Table 3
The conditional decomposition of the OLS estimation of the union density wage gap for total monthly wages

Variables	Base model (γ_0)	Full model (γ_1)	Gelbach decomposition
Union density wage gap	0.1619***	−0.0029	–
	(0.0244)	(0.0086)	
τ_{α_1} (Worker FE)	–	–	−0.0026
			(0.0055)
τ_{θ_1} (Job-title FE)	–	–	0.0571***
			(0.0198)
τ_{λ_1} (Firm FE)	–	–	0.1103***
			(0.0134)
R^2	0.5373	0.8801	

Notes: Decomposition based on Gelbach (2016). The base model includes as regressors a quadratic term in age, a quadratic term in tenure, schooling dummies (7), a gender dummy, firm size dummies (5), and sector dummies (25). The number of observations is 36,616,379. Robust firm clustered standard errors are in parentheses: *** denotes statistical significance at the 0.01 level. Sources: Quadros de Pessoal 1986-2009; Relatório Único, 2010-2013.

for the observable characteristics of the worker and the firm, the constant unobserved characteristics of workers, and the process of job-title placement, the worker compensation policies of firms are responsible for 11 log points of the union density wage gap of some 16.2 log points. Put differently, if every worker was faced with a neutral stance of his or her firm regarding its compensation policies, the union density wage gap would be reduced by about 68 percent. Next, consistent with the evidence provided in Panel A of Fig. 6, we find that the allocation of workers into job titles — either directly, or indirectly through promotion decisions — contributes 5.7 log points (or another 35 percent) of the estimated union density wage gap. Implicitly, therefore, trade unions achieve real success in either creating or in placing their members into higher paying job categories.[17]

Turning to the worker dimension, and after accounting for observable traits as well

variable is successively replaced by the corresponding estimated fixed effects. By construction, this decomposition is both exact (being the sum of the contributions corresponding to the difference between the two union density wage gap estimates) and unambiguous.

[17] In Table A.5 we use the ESCO occupation definition instead of job-title. Not surprisingly, the contribu-

as their sorting into firms and job-titles, it can be seen that individuals working in a fully-unionized firm receive compensation for their permanent unobserved characteristics that is estimated to be just 0.3 log points lower than in the case of a non-unionized firm. But to all intents and purposes there are no statistically significant differences between unionized and non-unionized workers in terms of (unobserved) ability.

In the pooled OLS estimation, the sources of identification of the union wage effect have both a cross-section and a time-varying dimension. The cross-section or between-effect arises simply from the comparison of wages of workers with different average union density rates. In the following decomposition exercise, we will explore the three sources of this cross-sectional variation: worker heterogeneity, firm heterogeneity, and job-title heterogeneity. This objective is achieved via the inclusion of their corresponding fixed (hence, cross-sectional) effects. The time-varying effects stem from two obvious sources: workers moving across firms with different union densities, and workers remaining in firms with changing union densities. The inclusion of a worker fixed effect necessarily removes the cross-sectional variation. The inclusion of the firm fixed effect precludes any variation arising from worker movement to other firms. The inclusion of the job-title fixed effect will restrict the source of wage variation to that which is not accounted by job-title moves. Thus, in our full model, identification of the union density wage gap comes solely from the wage effects of changing union density rates, net of (i.e. after accounting for) job-title changes. Not surprisingly, therefore, the union density wage gap is close to zero in the full model. The advantage of the full model, however, is that it allows us to obtain an unambiguous decomposition of the sources of heterogeneity and, by implication, to measure the relative importance of unobserved skills and sorting across firms and job-titles as drivers of the union wage premium.

Estimates of the Gelbach decomposition of the union density wage gap curve obtained from the kernel regression (equations (3) and (11)) are depicted in Fig. 7. The figure broadly confirms the principal result of the linear approach, namely the leading roles reserved for job titles and the compensation policies of firms (viz. the job-title and firm fixed effects). The flexibility of this approach indicates that the major source of non-linearity stems from the job-title component. The figure also provides a more informative picture of the sources of the union density wage gap. In particular, the assignment of workers to job titles reveals that there is a zone or a relevant region of union

tion of occupation fixed effects is much attenuated (to just 1.6 log points), suggesting that the informational content of the job-title fixed effects is more relevant in associating union density with wages.

densities (roughly between 50 and 85 percent) where its contribution to the wage gap is elevated, reaching almost 12.5 log points. Worker unobserved heterogeneity once again plays no role, irrespective of union density. Note, finally, that the part of the union density wage gap remaining after having allowed for the three high-dimensional fixed effects, and identified in the figure as the 'within component,' is essentially zero up to 70 percent density after which it increases to only a little over 2 log points at 100 percent density.

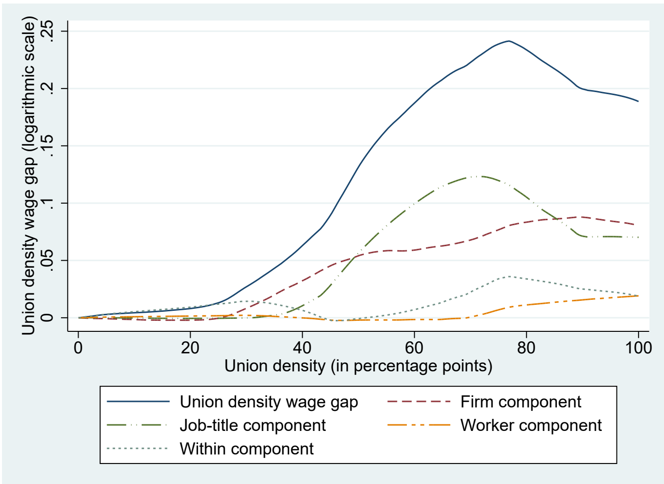

Fig. 7. The Gelbach decomposition of the union density wage gap curve for total monthly wages. Sources: Quadros de Pessoal, 1986-2009; Relatório Único, 2010-2013.

5.3. On the endogeneity of union membership

The main purpose of this study is to investigate what lies behind the union density wage gap rather than claiming a causal interpretation of the union density regression coefficient. Nevertheless, a number of endogeneity issues do need to be addressed. A widespread concern is that of individual self-selection into union status (Lee, 1978; Card et al., 2020). It may well be the case that more skilled workers disproportionately select into unionization, and in so doing drive the union density wage gap. This concern can be directly assessed within the framework of our decomposition exercise. Thus, there is no indication that union members have unobserved characteristics that render them more productive, given the very weak (near null) impact of the worker fixed effect on the wage gap.

Another potential source of ambiguity is the possibility that trade unions tend to organize in industries and firms where economic rents are higher. In this case, we face a problem of reverse causality. Although we certainly cannot rule out this possibility, we can nonetheless determine an upper-bound for this source of endogeneity. Specifically, the contribution of the firm fixed effect to the union density wage gap — of around two-thirds — is our upper-bound for the reverse effect. In other words, the indication that higher union density rates coincide with more generous firm wage policies may arise from 'pure' causal effects (unions pushing for higher wages) or from the tendency of trade unions to be organizationally stronger in those activities where economic rents are higher.

According to our methodology, the remaining one third of the union density wage gap is driven by the finding that union members are more likely to occupy better-paid job titles. This effect corresponds to the impact of job-title fixed effects on the wage gap. Because the wage floors for the job titles (i.e. bargained wages) and promotion policies are negotiated by the trade unions, this channel of wage determination is inherently linked to the actions and power of unions. In this sense, we would argue that the association between higher union density rates and better paid job titles can convincingly be regarded as causal. With some caution, therefore, we summarize this discussion by suggesting that the causal effect of union density on wages is likely to be contained within a 6 to 16 log points range.

Taking this source of causality one step further, one can think of the job-title fixed effect(s) as a suitable instrumental variable for union density. Using this set-up, the first-stage regression for union density will include the covariates of the full model. In the second stage, the coefficient estimate of predicted union density is exactly identified by the presence of job title in the first stage. On the basis of this alternative IV approach at least, it transpires that the two-stage coefficient estimate of the union density variable is a sturdier 0.126, or 12.6 log points (see Table A.6).

6. Inside the mechanisms of wage setting

In a complementary exercise, so as to better understand the role of trade unions in Portugal, it is instructive to decompose total wages into two parts: the bargained wage and a wage mark-up (Π). So the total wage is given by:

$$W_{Total} = W_{Bargained} \times \underbrace{\left[\frac{W_{Bargained} + \Pi}{W_{Bargained}}\right]}_{=1+\mu}, \tag{12}$$

where μ corresponds to the mark-up in relative terms. Through a straightforward logarithmic transformation, one obtains:

$$\ln W_{Total} = \ln W_{Bargained} + \underbrace{\ln(1+\mu)}_{\text{Wage cushion}} . \tag{13}$$

The bargained wage is the base wage floor as defined in the relevant collective agreement for the worker job title. For its part, the wage cushion corresponds to the difference between the total compensation and the bargained wage. As information on bargained wages is not contained in the dataset, we follow the methodology proposed in Cardoso and Portugal (2005), and define the bargained wage as the mode of the actual base wage within each year and job-title.[18] The wage cushion has two components. The first is simply the difference between the actual base wage and the bargained wage, as firms often pay a base wage above the bargained wage. The second component is the sum of the wage supplements received by the worker, to include seniority-indexed pay, overtime, fringe benefits such as the daily meals allowance, and irregularly paid components.

Table 4
Estimation of the union density wage gap for total monthly wages, the bargained wage, and the wage cushion

Dependent variable	Base model	R^2
Total compensation	0.1619	0.5373
	(0.0244)	
Bargained wage	0.2836	0.5095
	(0.0283)	
Wage cushion	−0.1218	0.0994
	(0.0184)	

Notes: The base model includes as regressors a quadratic term in age, a quadratic term in tenure, schooling dummies (7), a gender dummy, firm size dummies (5), and sector dummies (25). The number of observations is 36,616,379. Robust firm clustered standard errors are in parentheses: all coefficients are statistically significance at 0.01 confidence level. Sources: Quadros de Pessoal, 1986-2009; Relatório Único, 2010-2013.

[18] Having documented contractual wages in three industries employing around 10 percent of full-time workers in manufacturing and services, these authors show that the mode of the wage distribution of the base wage for each worker category within each collective agreement matches quite well the mandatory floors for each job-title at collective bargaining level.

We shall estimate two separate wage regressions, one for the (log) bargained wage and the other for the (log) wage cushion, where the wage cushion is expressed in relative terms. A useful way to look at the bargained wage regression is to think of an artificial exercise in which all workers collect the enacted wage floor, corresponding to their job titles, as signed in the applicable collective bargaining, and no more. On this basis, and as shown in Table 4, the union density wage gap would amount to 28.36 log points.

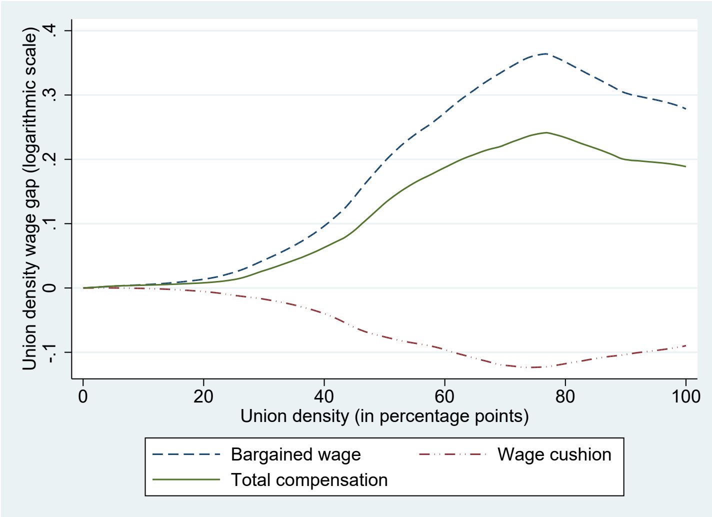

Fig. 8. The union density wage gap curve for total monthly wages, the bargained wage, and the wage cushion. Source: Relatório Único, 2010-2013.

At first glance, this may seem a puzzling result, as the union density wage gap for total monthly wages is only 16.19 log points, while that for the wage cushion is negative at -12.18 log points. Similar results are obtained when union density wage gap curves are estimated, as demonstrated in Fig. 8. Here the union density wage gap curve for the bargained wage attains levels of more than 30 log points when the large majority of the workforce is unionized, while the corresponding curve of the wage cushion declines up to around -12 log points. What is going on here is that the union density wage gap for total compensation hides the opposing effects of the two wage components.

The sizable union density wage gap of the bargained wage is not surprising, as unions naturally seek to lock in a significant share of their gains through the mandatory dispositions of collective agreements. Thus, where collective agreements are signed in sectors with highly unionized firms, signifying enhanced union bargaining power, unions may be expected to succeed in securing higher base wage floors (bargained wages). Our results in the form of a negative union density wage gap for the wage cushion do indeed suggest that the cushion is deployed by firms to attenuate the bargained wage gap. It follows that the union density wage gap for total compensation is lower than that for the bargained wage. This compression may result from either lower wage supplements or by smaller drift between the actual base wage and the bargained wage floor. In other words, in high union density environments, union success in raising the bargained wage limits the ability of firms to pay base wages in excess of bargained wages. In branches where trade unions are weaker, however, relatively low bargained wages offer scope for local improvement.

7. A glimpse into the impact of unions on wage inequality

We now take advantage of conditional quantile regression methods to explore the effect of union density on the entire wage distribution. The usefulness of conditional quantile regression to study the impact of unionization was first demonstrated by Chamberlain (1994) in his influential address to the 1990 World Congress of the Econometric Society. To the best of our knowledge, only the method of moments quantile regression estimator (MM-QR), recently proposed by Machado and Santos Silva (2019), is able to accommodate multiple high-dimensional fixed effects that can affect the whole distribution instead of just a location shift. Accordingly, we will use this estimator to analyze the impact of union density on wage inequality. A critical advantage of this approach is that it is based on the OLS solutions to two high-dimensional fixed effects regression equations, one corresponding to the location function and the other to the scale function. Apart from its ease of implementation, the MM-QR estimator lends itself in a straightforward way to the decomposition of the bias arising from the omission of the high-dimensional effects.

We will first specify the conditional quantile regression version for a model without fixed effects:

$$Q_w(\tau|(x_{it}) = \alpha^l + \boldsymbol{x'_{it}}\boldsymbol{\beta^l} + \sigma(\alpha^s + \boldsymbol{z'_{it}}\boldsymbol{\beta}^s)q(\tau), \tag{14}$$

where $Q_w(\tau|(x_{it})$ denotes the conditional quantile of w corresponding to percentile τ, z is a vector of transformed covariates x, and σ identifies the scale function. α^l and $\boldsymbol{\beta^l}$, and α^s and $\boldsymbol{\beta}^s$, correspond to the intercept and regression coefficients in the location and scale functions, respectively. $q(\tau)$ is simply $F^{-1}(\tau)$.

The full model can be written:

$$Q_w(\tau|(x_{it}) = \alpha_i^l + \lambda_f^l + \theta_j^l + \boldsymbol{x'_{it}\beta^l} + \sigma(\alpha_i^s + \lambda_f^s + \theta_j^s + \boldsymbol{z'_{it}\beta^s})q(\tau), \tag{15}$$

where α_i^l (α_i^s), λ_f^l (λ_f^s), and θ_j^l (θ_j^l) are the worker, firm, and job-title fixed effects in the location (scale) function, respectively.

Considering that $\boldsymbol{z = x}$ and that $\sigma(.)$ is the identity function, the model simplifies to:

$$Q_w(\tau|(x_{it}) = [\alpha_i^l + \alpha_i^s q(\tau)] + [\lambda_f^l + \lambda_f^s q(\tau)] + [\theta_j^l + \theta_j^s q(\tau] + \boldsymbol{x'_{it}[\beta^l + \beta^s} q(\tau)]. \tag{16}$$

Estimation of this model can be achieved in four simple steps.

(1) Obtain the parameters of the location function minimizing the sum of the squares of the residuals from the model:

$$w = \alpha_i^l + \lambda_f^l + \theta_j^l + \boldsymbol{x'_{it}\beta^l} + \epsilon^l, \tag{17}$$

using, for example, the algorithm of Guimarães and Portugal (2010). This procedure corresponds exactly to the OLS solution which was estimated before;

(2) Compute the estimation residuals from this model, $\hat{R_{it}}$, and calculate $|\hat{R}_{it}|$;

(3) Obtain the parameters of the scale function minimizing the sum of the squares of the residuals from the model:

$$|\hat{R}_{it}| = \alpha_i^s + \lambda_f^s + \theta_j^s + \boldsymbol{x'_{it}\beta^s} + \epsilon^s; \text{ and,} \tag{18}$$

(4) Obtain $q(\tau)$ as the τ-th sample quantile from a standardized residual $\hat{R_{it}}/(|\hat{R}_{it}|)$.

Results for the location-scale representation of the quantile regression model are given in Table 5. The first line of the table presents the union coefficient estimates corresponding to the base specification for the 10th, 50th, and the 90th percentiles.[19] The impact of union density at the conditional mean and at the conditional median are virtually identical (16.2 log points for the mean and 16.3 log points for the median). Low-wage workers (say those in the 10th percentile) record a higher union density wage gap (17.2 log points). For their part, high-wage workers (those in the 90th percentile) have a somewhat lower wage gap (15.0 log points) than the median wage worker. It

[19]The Stata procedure *mmqreg* provided by Santos Silva was used to obtain the method of moments quantile regression estimates.

Table 5
Method of moments quantile regression estimation

Variables	Percentiles		
	10	50	90
Base model			
γ_τ (Union density wage gap)	0.1720	0.1632	0.1498
	(0.0012)	(0.0012)	(0.0034)
Full model			
γ_τ (Union density wage gap)	-0.0052	-0.0031	-0.0003
	(0.1941)	(0.0996)	(0.1288)
Decomposition			
Worker FE	0.0001	-0.0024	-0.0056
Job-title FE	0.0710	0.0610	0.0476
Firm FE	0.1136	0.1068	0.1006
Residual	-0.0076	0.0009	0.0004

Notes: The base model includes as regressors a quadratic term in age, a quadratic term in tenure, schooling dummies (10), a gender dummy, firm size dummies (5), and sector dummies (25). All regression coefficients in the base model are statistically significant at the 0.01 confidence level. The regression coefficients in the full model are not statistically different from zero. The contributions of the fixed effects add the location and the scale the components. The number of observations is 36,616,379. Sources: Quadros de Pessoal 1986-2009; Relatório Único, 2010-2013.

appears that union membership mildly compresses the wage distribution by improving the wages of low-wage workers more than those of high-wage workers.

The full model specification, in which we take into account unobserved worker heterogeneity and sorting across firms with distinct wage policies and job titles is given in the second line of the table. Consistent with the previous linear regression model results is the suggestion that the union density wage gap is close to zero across all three percentiles. Nevertheless, worker, firm, and job-title heterogeneity fully account for the union density wage gap in the base model. Moreover, the three high-dimensional fixed effects account for the dispersion in the union coefficient estimates. In the same vein, we can extend the Gelbach procedure to decompose the change in the quantile regression coefficient estimates (from the base model to the full model). It is clear that we can apply the Gelbach decomposition to both the location and the scale functions. For the location function, the procedure is exactly the same (and, in turn, produces exactly the same results) as that followed in section 5.2. For the scale function, we obtain the contribution of each fixed effect again using the omitted variable bias formula and multiply it by $q(\tau)$ in order to differentiate them by percentile. There is, of course, a third (residual) component arising from the nonlinear change in $|\hat{R}_{it}|$ when we move from the base model to the full model.

Three indications emerge from this decomposition exercise. First, the contribution of of high-dimensional fixed effects is largely driven by the location function, which can be seen by comparing Table 5 with Table 3. Second, the wage compressing effect of union density is concentrated in the job-title dimension. Low-wage workers benefit more from being placed (or promoted) into higher paying job titles than do high-wage workers. Finally, the residual component of the decomposition is small, or even negligible, at higher percentiles.

8. Conclusions

This paper has shown that in a regime of near-universal collective bargaining coverage one may nevertheless discern pronounced union density wage gaps due to the heterogeneous influence of unions in covered settings. Our baseline estimates of the union density wage gap for total monthly wages top out at approximately 24 log points. Wage gaps are not proportional to union density. Rather, some critical mass (circa 30 percent density) is required to materially influence wages, while beyond some level (70 percent) further increases in union density detract from the peak premium.

Having obtained these sizable values of union density wage gaps based on estimates of average wage differentials between the wages of two observationally identical workers in two observationally identical firms with distinct levels of unionization, we turned to consider the potential sources of the gaps. Our fixed-effects model included in addition to worker and firm fixed effects a job-title or finely drawn occupational fixed-effect. We then applied Gelbach's (2016) decomposition which allows for the quantification of the share of the union density wage gap accounted for by each of these three fixed effects. That is, we sought to account for the separate contributions to the union density wage gap of the targeted wage policies of firms (i.e. versus a neutral stance which would imply that unionized workers are randomly assigned to firms); of differential rules, both direct and indirect, whereby unionized workers are assigned to higher paying job titles in the compensation tables of collective agreements; and of the manner of the allocation of workers of different unobserved ability to firms. The wage policies of firms were found to account for approximately two-thirds of the union density wage gap in monthly wages, a profound change over a situation in which workers are randomly assigned to firms. Also of importance was the allocation of unionized workers into better paying job titles. This source contributed another one-third of the union density wage gap. Accordingly, compensation for unobserved ability played little or no role in determining the premium enjoyed by union workers even if, as was demonstrated, their observed heterogeneity would patently produce a higher wage gap were it not controlled for.

We also distinguished between the components of total earnings to gain a fuller understanding of the role of trade unions. Specifically, we split total earnings into the bargained wage and the wage cushion. The bargained wage establishes the rate for the highly detailed job classifications or job titles that are contained in each collective bargaining agreement and sets a floor to the base wage obtaining at firm level. For its part, the wage cushion is the gap between the bargained wage and the actual base wage determined at firm level plus the difference between the latter and total earnings. We found that the size of the union density wage gap is considerably larger for the bargained wage than for total earnings — peaking at more than 30 log points once some three-quarters of the workforce is unionized — and is actually negative for the wage cushion given its specification as the logarithmic ratio of actual earnings to the bargained wage. In other words, the union density wage gap would be much higher (and wage dispersion lower) were all workers remunerated according to the bargained wage. The disparity between the premia attaching to the bargained wage and total earnings, and the negative gap for the wage cushion, reflect firm-specific arrangements that partly offset collective bargaining. In those sectors where trade unions are weaker, the relatively low bargained wages offer scope for local adjustments as when employers experience market opportunities or face constraints.

In a final contribution we sought to determine union impact on wage inequality, using conditional quantile regressions. Here, we used an approach that can accommodate multiple high-dimensional fixed effects, using the methods of moments quantile regression estimator to explore the impact of union density on the whole wage distribution. Estimation of the baseline model suggested that union membership mildly narrowed the wage distribution. For the full model containing the three fixed effects, application of the Gelbach procedure indicated that this narrowing effect of union density was concentrated in the job-title dimension where low wage workers benefited more from being placed or promoted into higher-paying job titles than other groups.

In sum, the present treatment can report progress along several dimensions. It uses an innovative procedure to determine the union density premium in circumstances, again by no means unique to Portugal, where almost all workers are covered by a collective agreement but few are union members. It attributes the union density wage gap to three types of heterogeneity that ‘explain’ almost all wage variation. These high-dimensional fixed effects were also directly incorporated in a new application investigating the effect of union density on the wage distribution. They also contributed to our discussion of endogeneity, leading us to conclude that the causal effect of union density on wages likely fell within a 6 to 16 log points range. For the future, more work is clearly needed on the firm fixed effect. Reflecting recent developments in the literature, a logical extension would be to examine the contribution of firm monopsony power and the countervailing influence of union density.

Acknowledgements

The authors benefited from the comments of Lars Vilhuber (guest editor), three referees, session participants at the 2017 ASSA Annual Meeting in Chicago, the 2019 Models of Linked Employer-Employee Data Conference, and those of seminar discussants at the CEP-LSE and the Bank of Portugal. They also wish to thank Ana Rute Cardoso, Henry Farber, Alan Manning, Stephen Machin, Pedro Martins, Steve Pischke, and John Van Reenen for their helpful remarks. Lucena Vieira provided outstanding computational assistance. Addison is indebted to the Riegel and Emory Human Resources Center at the University of South Carolina for generous research support. Portugal gratefully acknowledges financial support from Fundação para a Ciência e Tecnologia under grant PTDC/EGE-ECO/7493/2020. De Almeida Vilares gratefully acknowledges financial support from the Economic and Social Research Council, grant #1642247, and from Fundação para a Ciência e Tecnologia, grant #2020.04812.BD.

References

Abowd, John M., Francis Kramarz and David N. Margolis, 1999. 'High Wage Workers and High Wage Firms', *Econometrica* **67**(2), 251–333.

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.

Blanchflower, David and Alex Bryson, 2003. 'Changes Over Time in Union Relative Wage Effects in the UK and the USA Revisited', In Addison John T., Claus Schnabel (eds.), *International Handbook of Trade Unions,* pp. 197–245, Cheltenham, UK: Edward Elgar Publishing Ltd.

Card, David and Ana Rute Cardoso, 2021. 'Wage Flexibility Under Sectoral Bargaining', NBER Working Paper No. 28695, National Bureau of Economic Research, Cambridge MA.

Card, David, Thomas Lemieux and W. Craig Riddell, 2004. 'Unions and Wage Inequality', *Journal of Labor Research* **25**(4), 519–562.

Card, David, Thomas Lemieux and W. Craig Riddell, 2020. 'Unions and Wage Inequality: The Roles of Gender, Skill and Public Sector Employment', *Canadian Journal of Economics* **53**(1), 140–173.

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 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.

Chamberlain, Gary, 1994. 'Quantile Regression, Censoring, and the Structure of Wages', In Sims, Christopher. A. (ed.), *Advances in Econometrics (Sixth World Congress)* **Vol. 1**, 171–210, Cambridge University Press.

DiNardo, John, Nicole Fortin and Thomas Lemieux, 1996. 'Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach', *Econometrica* **64**(5), 1001–44.

Farber, Henry S, 1986. 'The Analysis of Union Behavior', *Handbook of Labor Economics* **2**, 1039–1089.

Farber, Henry S., Daniel Herbst, Ilyana Kuziemko and Suresh Naidu, 2021. 'Unions and Inequality Over the Twentieth Century: New Evidence from Survey Data', NBER Working Paper No. 24587, National Bureau of Economic Research, Cambridge, MA.

Frandsen, Brigham R, 2012. 'Why Unions Still Matter: The Effects of Unionization on the Distribution of Employee Earnings', Mimeograph, Massachusetts Institute of Technology, Cambridge, MA .

Gelbach, Jonah B, 2016. 'When Do Covariates Matter? And Which Ones, and How Much?', *Journal of Labor Economics* **34**(2), 509–543.

Guimarães, Paulo and Pedro Portugal, 2010. 'A Simple Feasible Procedure to Fit Models with High-dimensional Fixed Effects', *Stata Journal* **10**(4), 628–649.

Hijzen, Alexander and Pedro S Martins, 2016. 'No Extension Without Representation? Evidence from a Natural Experiment in Collective Bargaining', IMF Working Paper WP/16/143, Research Department, International Monetary Fund, Washington, DC .

Hirsch, Barry T, 2004. 'Reconsidering Union Wage Effects: Surveying New Evidence on an Old Topic', *Journal of Labor Research* **25**(2), 233–266.

Lee, David S and Alexandre Mas, 2012. 'Long-Run Impacts of Unions on Firms: New Evidence from Financial Markets, 1961–1999', *The Quarterly Journal of Economics* **127**(1), 333–378.

Lee, Lung-Fei, 1978. 'Unionism and Wage Rates: A Simultaneous Equations Model with Qualitative and Limited Dependent Variables', *International Economic Review* **19**(2), 415–33.

Machado, José A. F. and J. M. C. Santos Silva, 2019. 'Quantiles Via Moments', *Journal of Econometrics* **213**(1), 145–173.

Silverman, Bernard W, 1986. *Density Estimation for Statistics and Data Analysis*, London: Chapman and Hall.

Appendix

Table A.1
The most representative collective agreements, 2010-2013

Rank	Collective agreement description	Number of covered workers per year	Average union density	Number of job titles
1	Building and construction	254,840	2.9%	461
2	Office workers	143,941	2.9%	62
3	Metalwork and metallurgy industry	126,300	10.4%	34
4	Building and construction II	109,115	2.6%	421
5	Supermarkets and hypermarkets	108,632	9.4%	125
6	Private social security institutions	103,563	5.8%	386
7	Car Industry and trade	78,897	7.7%	412
8	Building and construction III	69,662	2.2%	407
9	Textile industry	65,065	5.6%	233
10	Hospitality sector - Center and South region - restaurants	63,072	3.0%	259
11	Retail - Lisbon region	61,141	2.9%	493
12	Cleaning services	60,240	8.7%	105
13	Hospitality sector - North region - restaurants	45,364	1.6%	390
14	Public transport and road transport of goods	43,470	3.4%	106
15	Private education	42,196	3.2%	115

Source: Relatório Único, 2010-2013.

Table A.2
OLS estimation of the union density wage gap for total monthly wages including non-covered workers

Variable	Coefficient (s.e.)	
Union density	0.1370	0.1387
	(0.0197)	(0.0197)
Non-covered worker		0.0605
		(0.0088)
R^2	0.5499	0.5503

Notes: The dependent variable is total monthly wages (in logs). The controls also include a quadratic term in age, a quadratic term in tenure, a gender dummy, schooling dummies (7), firm size dummies (5), sector of activity dummies (25), and year dummies (24). The number of observations is 43,455,566. Robust firm clustered standard errors are in parentheses: all coefficients are statistically significant at the 0.01 confidence level. Sources: Quadros de Pessoal 1986-2009; Relatório Único, 2010-2013.

Table A.3
Summary statistics

Variable	Period	
	1986-2009	2010-2013
Union density	–	9.90%
	–	(21.77)
Logarithm of total monthly wages	5.47	5.60
	(0.55)	(0.53)
Logarithm of bargained wages	5.22	5.31
	(0.46)	(0.44)
Logarithm of wage cushion	0.25	0.29
	(0.40)	(0.43)
Worker's age (in years)	36.85	40.01
	(10.93)	(10.37)
Tenure of the worker (in years)	8.18	8.79
	(8.43)	(8.56)
Proportion of females	40.4%	45.7%
Education:		
less than primary school	2.76%	0.68%
primary school	35.59%	15.52%
basic school	21.94%	18.51%
elementary school	17.05%	27.05%
secondary school	16.39%	24.2%
post-secondary school	0.23%	0.41%
University attendance	1.58%	1.90%
College degree	4.47%	11.72%
Workers in firms with:		
less than 50 employees	47.23%	51.94%
50 to 99 employees	11.73%	11.51%
100 to 499 employees	21.20%	18.92%
500 to 999 employees	6.00%	4.55%
1000 to 4999 employees	8.93%	8.13%
more than 5000 employees	4.91%	4.95%
Worker-year observations	30, 397, 602	6, 237, 187

Notes: For the continuous variables, the mean is presented, with the standard deviation displayed in parenthesis. Sources: Quadros de Pessoal, 1986-2009; Relatório Único, 2010-2013.

Table A.4
Conditional decomposition of the OLS estimation of the union density wage gap for different model specifications

Variables	Base model (γ_0)	Full model (γ_1)	Gelbach decomposition
Standard model			
Union density wage gap	0.1619***	−0.0029	–
	(0.0244)	(0.0086)	
τ_{α_1} (Worker FE)	–	–	−0.0026
			(0.0055)
τ_{θ_1} (Job-title FE)	–	–	0.0571***
			(0.0198)
τ_{λ_1} (Firm FE)	–	–	0.1103***
			(0.0134)
R^2	0.5373	0.8801	
Standard model without firm size and sector dummies.			
Union density wage gap	0.4492***	−0.0222*	–
	(0.0333)	(0.0115)	
τ_{α_1} (Worker FE)	–	–	0.0031
			(0.0063)
τ_{θ_1} (Job-title FE)	–	–	0.0637**
			(0.0281)
τ_{λ_1} (Firm FE)	–	–	0.4047***
			(0.0208)
R^2	0.4442	0.8799	
Standard model without education, firm size and sector dummies.			
Union density wage gap	0.6643***	−0.0178	–
	(0.0572)	(0.0123)	
τ_{α_1} (Worker FE)	–	–	0.0744***
			(0.0140)
τ_{θ_1} (Job-title FE)	–	–	0.1491***
			(0.0337)
τ_{λ_1} (Firm FE)	–	–	0.4586***
			(0.0261)
R^2	0.2044	0.8797	

Notes: Decomposition based on Gelbach (2016). The standard model includes as regressors a quadratic term in age, a quadratic term in tenure, schooling dummies (7), a gender dummy, firm size dummies (5), year dummies (24), and sector dummies (25). Robust firm clustered standard errors are in parentheses: ***, **, * denote statistical significance at the 0.01, 0.05, and 0.10 levels, respectively. The number of observations is 36,616,379. Sources: Quadros de Pessoal, 1986-2009; Relatório Único, 2010-2013.

Table A.5
The conditional decomposition of the OLS estimation of the union density wage gap for total monthly wages, using ESCO occupation definition

Variables	Base model (γ_0)	Full model (γ_1)	Gelbach decomposition
Union density wage gap	0.1619***	0.0091	–
	(0.0244)	(0.0095)	
τ_{α_1} (Worker FE)	–	–	0.0072
			(0.0086)
τ_{θ_1} (Occupation FE)	–	–	0.0164**
			(0.0077)
τ_{λ_1} (Firm FE)	–	–	0.1292***
			(0.0126)
R^2	0.5373	0.8801	

Notes: Decomposition based on Gelbach (2016). The base model includes as regressors a quadratic term in age, a quadratic term in tenure, schooling dummies (7), a gender dummy, firm size dummies (5), and sector dummies (25). Robust firm clustered standard errors are in parentheses: ***, ** denote statistical significance at the 0.01 and 0.05 levels, respectively. The number of observations is 35,764,082. Sources: Quadros de Pessoal 1986-2009; Relatório Único, 2010-2013.

Table A.6
2SLS estimation of the union density wage gap for total monthly wages, using the job-title fixed effects as an instrument for union density

Variable	2SLS
Union density	0.1257
	(0.0658)
Worker FE	Yes
Firm FE	Yes
R^2	0.8683

Notes: The dependent variable is total monthly wages (in logs). The second stage regression also includes as controls a quadratic term in age, a quadratic term in tenure, a gender dummy, schooling dummies (7), firm size dummies (5), sector dummies (25), year dummies (24), firm fixed effect and worker fixed effect. The first stage regression has union density on the left-hand side and the job-title fixed effect plus the controls presented in the second stage on the right-hand side. Robust firm clustered standard error is in parentheses. The coefficient for predicted union density is statistically significant at the 0.10 confidence level. The number of observations is 36,616,379. Sources: Quadros de Pessoal 1986-2009; Relatório Único, 2010-2013.