Three Essays on the Economics of Inequalities and Discrimination

(Article begins on next page)

Università degli Studi di Torino

Dipartimento di Scienze Economico-Sociali e Matematico-Statistiche

"V. Pareto" Doctoral Program in Economics

XXIX Cycle

S.S.D.: SECS-P/01

Three Essays on the Economics of Inequalities and Discrimination

Bernardo Fanfani

Supervisor: prof. Francesco Devicienti

A thesis presented for the degree of Doctor of Philosophy

2014 - 2017

Table of Contents

Chapter I
Tastes for Discrimination in Monopsonistic Labour Markets

Chapter II
Collective Bargaining and the Evolution of Wage Inequality in Italy

Chapter III
The Gender Wage Gap Among Italian Employees. Evidences from the ISFOL PLUS Database

Chapter I

Tastes for Discrimination in Monopsonistic Labour Markets*

Abstract

This paper studies a model of the gender wage gap where differences between men and women arise from two main mechanisms: taste-based discrimination and monopsonistic power. The theoretical results inform an empirical analysis of the gender wage gap in the Italian manufacturing sector, for which we propose a novel methodology that, under reasonable assumptions, allows the identification of firm-specific taste-based discrimination parameters. Our approach relies on the use of a matched employer-employee database and it exploits the relatively homogeneous labour market structure faced by firms belonging to the same local labour market (*i.e.* commuting area characterized by a high density of manufacturing firms). Using this method -and focusing mainly on the effects of management structure and of the female share of employment within firms- we test several theoretical implications of our model, which have not been fully addressed by the previous literature. Our results show that taste-based discrimination is a potentially important determinant of the overall gender wage gap. Moreover, the absence of female workers at the top of the firms' hierarchy, as well as the female share of employment within workplaces, represent two valid proxies of firm-specific preferences against women.

JEL Codes: J00, J16, J23, J3, J7.

Keywords: Gender Wage Gap; Taste-Based Discrimination; Monopsonistic Discrimination; Firm Wage Policy; Matched Employer-Employee Data.

*I would like to thank Ronald Oaxaca, David Card, Francesco Devicienti, Ignacio Monzon for their excellent comments and the invaluable help and support they have provided me with. I also would like to thank participants at Collegio Carlo Alberto workshops and UC Berkeley seminars. Part of this work has been written during a visiting period spent at the Department of Economics of UC Berkeley, which I would like to thank for the hospitality.

1 Introduction

Gender wage gaps are one of the most persistent economic regularities, on which many hypothesis have been formulated. In this paper, we follow an approach that combines elements of several of the existing theories.[1] In particular, we build a model where gender pay differences are determined by two main mechanisms: *Becker-type* (or so-called taste-based) and *Robinsonian* (or so-called monopsonistic) discrimination. Based on this model, we develop an empirical approach to identify, in the context of an imperfect labour market, the presence of taste-based discrimination. Moreover, we derive formal tests on the validity of proxies that are usually taken as firm-specific measures of discriminatory preferences. Finally, by applying these methods on Italian matched employer-employee data, we find empirical evidence supporting the validity for two commonly used proxies of taste-based discrimination, *i.e.* the presence of women in managerial positions and the share of female employment within firms.

According to the theory of Becker [1957], taste-based discrimination arises because some employers have a dis-utility in working with women, so that either they are able pay them less than their *fair share*, or they avoid hiring them, reducing the aggregate female labour demand. As a consequence, all employers with small enough discriminatory preferences are able to hire a given quantity of female workers at a lower wage than the one needed to hire the same quantity of men.[2] Instead, *Robinsonian* discrimination is a mechanism arising when employers have monopsonistic power in the factor market. When relaxing the assumption of price taking behaviour, employers minimize costs not only on the extensive margin, by adjusting quantities, but also on the intensive margin, by adjusting wages. In this context, according to the *Robinsonian* discrimination hypothesis, gender wage differences are at least in part driven by employers' greater monopsonistic wage-setting power against women, given that, on average, the female labour supply to the firm is more rigid

[1] See Blau and Kahn [2017] for a recent literature review on the main theories and existing evidences on the gender wage gap.

[2] See, among others, Charles and Guryan [2008] for a discussion and an evaluation of several implications of Becker's theory in the context of the racial wage gap.

than the male one.[3]

Assuming that labour markets are not perfectly competitive can be motivated on both, theoretical and empirical grounds. Several studies, using different approaches in a variety of contexts, have found evidences documenting come degree of employers' wage setting power and substantial differences in the female and male labour supply elasticities to the firm.[4] Even if some exceptions to this trend can also be found in the empirical literature, to the best of our knowledge there are no overwhelming evidences against the hypothesis that most labour markets are not perfectly competitive.[5]

The choice of a monopsonistic model is grounded also on more theoretical considerations. Boal and Ransom [1997] show that monopsonistic labour markets are implied by several dynamic search models in which larger firms face diseconomies of scale in hiring workers. On this respect, Black [1995] builds a dynamic model where taste-based discrimination itself produces monopsonistic discrimination against minority groups. Moreover, Boal and Ransom [1997] also show that there are many other potential mechanisms that would imply imperfect competition in the labour market, among which imperfect information about vacancies and costs associated to mobility from a job to the other.

In a recent contribution, Card et al. [2017] argue that firms' monopsonistic power could also represent an important mechanisms driving heterogeneity in wages between observationally similar workplaces, an evidence documented by several studies following the seminal work by Abowd et al. [1999]. This feature of the model is of particular interest for the purposes of the present analysis, especially given that another recent work by Card et al. [2016] provides novel evidence on the relevant role of workplace heterogeneity as a

[3]The original model of monopsony dates back to the 1930s (Robinson [1933]), but interest on the relationship between the labour market structure and gender discrimination has emerged only more recently (see Boal and Ransom [1997] and Manning [2003]).

[4]Among others, see Barth and Dale-Olsen [2009], Hirsch et al. [2010], Ransom and Sims [2010], Ransom and Oaxaca [2010], Depew and Srensen [2013], Muehlemann et al. [2013], Webber [2015] and, for Italy, Sulis [2011]. All of these studies provide either indirect or direct support to the hypothesis of monopsonistic labour markets.

[5]One notable exception supporting the perfectly competitive hypothesis is the paper by Matsudaira [2014]. This study exploits a quasi-experimental variation in the size of firms, *i.e.* the variation induced by a minimum staffing law introduced in California's caring sector. Results show that there was no growth in wage differences between nursing homes affected and not affected by this reform. However, since around 75% of all nursing homes in the State were under-staffed according to the new law, the resulting strong growth in aggregate demand for caregivers could have produced general equilibrium effects that would justify similar results also in monopsonistic labour markets.

determinant of the gender pay gap. Indeed, studying Portuguese data, these authors find that firm wage policies paid to women are consistently lower than those paid to men and they attribute this outcome to rent sharing mechanisms within firms, which seem to be somewhat more disadvantageous for women.

Our analysis contributes to this literature by showing that a simple static model of taste-based and monopsonistic discrimination provides a coherent theoretical framework for interpreting such firm-specific component of the gender wage gap. Moreover, the main contribution of this paper is to show that, within narrowly defined geographical and economic contexts, it is possible to identify and test the validity of potential proxy variables for employers' discriminatory preferences. Finally, we show how this approach allows to obtain fairly close approximations to what is the overall impact of taste-based discrimination on the gender wage gap. Even if this last result is undermined by a not very neat distinction between classical discrimination effects and a few possible confounding factors, most notably measurement errors and within firm gender differences in compensating wage differentials, our approach sheds light on a component of gender pay gap that is usually left unexplained by traditional cross-sectional methods.

Being able to distinguish among the sources of wage differences between men and women is not merely a theoretical exercise, but it has important implications on the choice of the most effective policies to implement in order to achieve greater equality. A second interesting feature of our analysis is that the method proposed here can also be applied to provide more solid ground to studies that aim at testing the implications of Becker's theory in the data. For example, a particularly important stream of literature concerns the relationship between taste-based discrimination and firms' product market structure. According to short-run predictions of this model, employers hiring more members of the disadvantaged group should have lower costs and be more profitable.[6] Moreover, in the long run gender differences should reduce, given that discriminatory firms are less efficient than incumbent non-discriminatory competitors.[7]

[6]This outcome is studied, among others, by Hellerstein et al. [2002] and Kawaguchi [2007]

[7]Studies on Becker's long-run prediction often exploit shocks in product market competition across time, but the magnitude and the extent of their effects on discrimination are often found to be particularly limited (*e.g.* Black and Brainerd [2004] and Heyman et al. [2013]). Among studies adopting

Virtually all of the existing contributions on this topic face the challenge of identifying a reliable parameter for discriminatory preferences. Such parameter is sometimes explicitly available from survey data only at an aggregate level (as for example in Charles and Guryan [2008]). More often, discriminatory preferences are approximated by the female share of workers within firms (*e.g.* Weber and Zulehner [2014]) or, in some cases, by the presence of women in executive boards (*e.g.* Flabbi et al. [2014]).[8] To the best of our knowledge, this paper is the first to propose an empirical test on the validity of firm-specific measures of taste-based discrimination.

The application of this paper is based on data covering the population of private sector workers in the Veneto region of Italy. We first measure the relative importance of firm-specific heterogeneity as a determinant of the overall gender pay gap, following the methodology of Card et al. [2016]. However, instead of estimating the impact of economy-wide workplace heterogeneity on the gender wage gap, we focus our analysis on manufacturing *local labour markets* only, since these can be considered groups of firms characterized by relatively homogeneous labour market structures.[9]

Given that the selected firms share a very similar labour supply, while market-wide gender differences in human capital and returns to skills are fully taken into account by the regression model, we are able to provide conditions under which given comparisons between firm-specific wage components provide a meaningful representation of discriminatory preferences. We then test several proxies of taste-based discrimination, finding strong support on the validity of presence of women in the management and female share of employment within firms as proxies of more or less discriminatory behaviour against women. Moreover, we test also whether other mechanisms are in place, finding some sup-

different approaches, Ashenfelter and Hannan [1986] find a negative relationship between women's employment shares and market concentration in the banking industry. More recently, Weber and Zulehner [2014] find that firms with larger female shares have better survival rates, while surviving firms tend to increase women's employment levels, an evidence supporting the hypothesis of employers' learning and of convergence toward less discriminatory outcomes.

[8] As we show in the paper, even from a purely theoretical point of view the female share of workers within firms is not a perfect proxy for taste-based discrimination if labour markets are monopsonistic.

[9] These local labour markets, alternatively called *districts*, are geographical and economic entities characterised by a high density of small-sized manufacturing-oriented firms. Such entities are defined by the Italian national statistical office using census data on commuting behaviour. For the purposes of our analysis, we apply an even stricter definition based on observed job mobility patterns.

port on the hypothesis that compensating wage differentials may be a second important driver of the gender wage gap in firms' pay policies.

The paper is organized as follows. Section 2 presents the theoretical model. Section 3 discusses the identification of firms' wage policies. Section 4 presents the data and discusses the identification of local labour markets. Section 5 presents the main empirical results of the paper, while the final section contains the concluding remarks.

2 Theoretical framework

2.1 Profit maximization

We consider a model where *Robinsonian discrimination* arises as the result of third degree price discrimination, while taste-based discrimination is defined as an employer-specific exogenous cost, which is proportional to the female employment level. In this model, an employer chooses a quantity of labour $L = L^m + L^f$ maximizing the profit function, which reads as

$$\pi(L^m, L^f) = pq(L) - w^m(L^m)L^m - w^f(L^f)L^f - \delta L^f \tag{1}$$

Throughout the paper, the subscripts m and f stands for male and female respectively. The parameter p is the output price and $q(L)$ is the quantity produced, with $q' > 0$ and $q'' < 0$. Male and female workers are perfect substitutes in the technology. w^m and w^f are gender-specific inverse labour supply functions, which are increasing in L^m and L^f, respectively. Finally, δ is a taste-based discrimination parameter.

Under standard assumptions,[10] the first order conditions of profit maximization can be written as

$$mp = w^m\left(1 + \frac{1}{\epsilon^m}\right) \qquad\qquad mp = w^f\left(1 + \frac{1}{\epsilon^f}\right) + \delta$$

where mp is the marginal revenue product and ϵ^g is the elasticity of the labour supply for $g = m, f$. The solution of the model is graphically represented in Figure 1, where the

[10] Two sufficient conditions for optimality are

$$2w^{g\prime} + w^{g\prime\prime}L^g > 0 \qquad\qquad \text{for } g = m,\ f$$

optimality conditions are characterized under different choices of the parameters. Namely, the left panel of the figure represents the solutions when $\epsilon^m \to \infty$ and $\epsilon^f < \infty$ for the cases of zero and positive taste-based discrimination, while the right panel describes the solutions when $\epsilon^g < \infty$ $(g = m, f)$, again for the cases in which $\delta = 0$ and $\delta > 0$.

In general, in this model wages are marked-down with respect to the marginal revenue product, and this mark-down grows as the labour supply becomes more rigid. Moreover, if $\delta = 0$ the marginal revenue product is set equal to each gender-specific marginal factor cost. When $\delta > 0$, an employer reduces the female employment and wage levels and this reduction is compensated by only a less than proportional growth in male employment, since hiring more men is increasingly costly, unless the male labour supply is perfectly elastic.

To sum up, an employer for which $\delta > 0$ produces less output, hires less women, has a lower female share and pays women less than what would be observed at the monopsonistic benchmark (*i.e.* at $\delta = 0$). However, when comparing any two firms, the negative relationships between any of these variables and employers' discriminatory preferences do not necessarily hold, since each firm may have different labour supply functions and production technologies. For this reason, in the next paragraph we characterize employers' heterogeneity more explicitly.

2.2 The Role of Workplace Heterogeneity

In this paragraph we characterize differences in average wages across firms in the context of the profit maximization model discussed above. We also introduce the possibility of heterogeneity in individual labour productivity, by allowing workers to provide different contributions to firms' revenues. Finally, we discuss the economic interpretation of firm wage residuals estimated in the context of an AKM regression model (see Abowd et al. [1999]), characterizing gender differences in such residuals.

Consider a population of firms indexed by j, each facing arbitrary gender-specific inverse labour supply functions. We assume that there is heterogeneity in productivity across workers, denoting with mp_j^i the marginal revenue product of employee i of a given gender

Figure 1: **Graphical Representation of the Model Under Different Market Strucutres and Discrimination Levels**

($\boldsymbol{a}$)
Labour Market Competitive for Men, Monopsonistic for Women

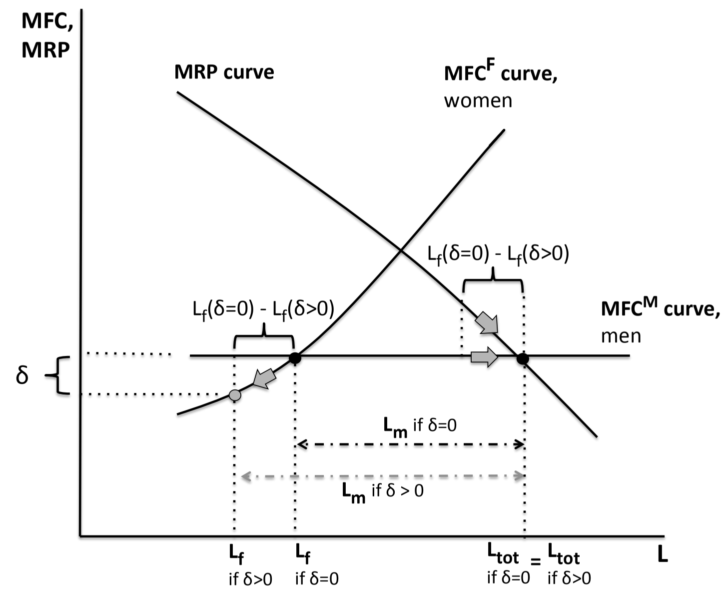

($\boldsymbol{b}$)
Monopsonistic Labour Market for Men and Women

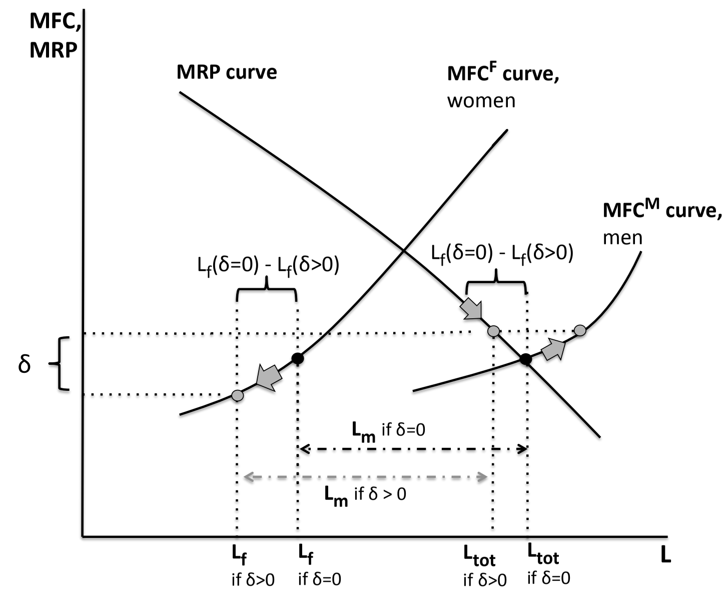

($\boldsymbol{a}$) The equilibrium in the absence of taste based discrimination ($\delta = 0$) is represented by the black dots, where $MRP = MFC^F\,[L_f(\delta = 0)] = MFC^M$. If $\delta > 0$, female employment is reduced to the level $L_f(\delta > 0)$, represented by the grey dot. Instead, MRP and MFC^M are kept constant, which implies that the difference $L_f(\delta > 0) - L_f(\delta = 0)$ is replaced with male workers. In this setting, for constant labour supply functions δ affects the share of female workers more, the smaller the overall firm size. Therefore, a large firm with $\delta = 0$ may have a lower share of women than a small taste-discriminatory firm.

($\boldsymbol{b}$) The equilibrium in the absence of taste based discrimination ($\delta = 0$) is represented by the black dots, where $MRP = MFC^F\,[L_f(\delta = 0)] = MFC^M\,[L_m(\delta = 0)]$. The grey dots represent the optimal points on each of these curves when $\delta > 0$. At this equilibrium, female employment is reduced to the level $L_f(\delta > 0)$, represented by the grey dot on the MFC^F line. As a consequence, MRP grows above the level $MFC^M\,[L_m(\delta = 0)]$. Therefore, more male workers are hired until MRP and $MFC^M\,[L_m(\delta > 0)]$ are equal again (grey dots on the respective lines). The growth in male employment is smaller than $L_f(\delta > 0) - L_f(\delta = 0)$ (*i.e.* for fixed technology and supply functions, more discriminatory firms are smaller in size). For fixed labour supply curves *and* conditional on total employment (L), the higher δ the lower the share of female workers.

g within firm j. That is to say, each worker provides a specific marginal contribution to firm's product revenues and this is quantity is known to employers. With this assumption, the first order condition of profit maximization becomes employee-specific and can be written as

$$w_j^i = mp_j^i \left(\frac{\epsilon_j^g}{1+\epsilon_j^g} \right) \left(1 - 1\left[g=f\right] \frac{\delta_j^i}{mp_j^i} \right)$$

Notice that in the above equation δ_j^i is modelled as an employer- *and employee*-specific discriminatory parameter for all female workers. However, we will now impose more structure on it and express employers' discriminatory preferences as a percentage (ψ_j) of female employees' productivity, assuming that *within firms* this proportion is constant across female workers of different quality.[11] More precisely, we define the following parameters

$$\begin{aligned} -\hat{\delta}_j &\equiv \ln\left(1-\delta_j\right) \\ \delta_j &\equiv \frac{\delta_j^i}{mp_j^i} = \psi_j \qquad\qquad \forall \left\{ i :\ i \in j \ \wedge \ i \in f \right\} \end{aligned}$$

where i is an individual identifier, while j and f are both firm (gender) identifiers and sets of individuals belonging to the corresponding workplace (gender group). δ_j^i is a parameter that is higher for more productive female workers within the firm. Finally, the parameter $\hat{\delta}_j$ is monotonic and increasing in δ_j, it is constant at the firm level and it describes the percentage of women's productivity that is marked-down due to employer's prejudices. Using the above definitions, firms' (geometric) average wages, denoted by w_j^m and w_j^f, can be written as

$$\begin{aligned} w_j^g &= mp_j^g \left(\frac{\epsilon_j^g}{1+\epsilon_j^g} \right) e^{-\hat{\delta}_j 1[g=f]} \qquad\qquad g = m, f \\ \implies \ln w_j^g &= \ln mp_j^g + \ln\left(\frac{\epsilon_j^g}{1+\epsilon_j^g} \right) - \hat{\delta}_j 1\left[g=f\right] \end{aligned}$$

[11]For example, if ψ_j is equal to 0.01 employer j would be indifferent between paying a male manager 10,000 Euro and a female manager 9,900 Euro, or similarly paying a male blue collar 1000 Euro and a female one 990 Euro. If we had modelled ψ_j as a lump-sum amount for each woman hired, the effects of taste-based discrimination would be relatively weaker for more productive workers, something probably less realistic given the vast evidence on *glass ceiling* effects (*e.g.* Arulamplam et al. [2006]).

where mp_j^g is the (geometric) average marginal revenue product of employees of gender g within firm j. In order to allow for the possibilities of measurement error and model misspecification, we now introduce additional components to the firms' wage equations. η_j is a residual representing firms' deviations from the labour supply schedules equally affecting men and women. Such deviations may be attributed to measurement error, efficiency wages, compensating wage differentials and employers' rent-sharing policies. ρ_j^m and ρ_j^f are gender specific residuals equally affecting male or female employees within a firm. These components may be interpreted partly as estimation errors, partly as gender differences in rent-sharing or wage incentives within the workplace. Finally, r_i^m and r_i^f are individual specific wage residuals, which we assume to be normally distributed with mean zero in the population and independent from all the other wage components. Adding these elements to the wage equation, we have a model that reads as

$$\ln w_j^g = \ln mp_j^g + \underbrace{\ln\left(\frac{\epsilon_j^g}{1+\epsilon_j^g}\right) - \hat{\delta}_j 1\left[g=f\right] + \eta_j + \rho_j^g}_{\equiv\, \omega_j^g} + \frac{1}{L_g}\sum_{i\in j} r_i^g \qquad (2)$$

Given the conditions imposed on r_i^g and assuming that workers' productivity can be approximated correctly by time-constant and time-varying individual characteristics, the element ω_j^g defined above can be interpreted as a time-constant firm wage residual, which can be recovered in the context of an AKM regression model estimated separately by gender.[12] Throughout the paper, we call this residual *firm wage policy*, or *firm wage premium.*

Taking differences between male and female firm pay policies, we have

$$\omega_j^m - \omega_j^f = \ln\left(\frac{\epsilon_j^m}{1+\epsilon_j^m}\right) - \ln\left(\frac{\epsilon_j^f}{1+\epsilon_j^f}\right) + \hat{\delta}_j + \underbrace{\rho_j^m - \rho_j^f}_{\equiv\, \rho_j} \qquad (3)$$

According to the above equation, gender differences in firms' wage policies are determined by a combination of: gender differences in monopsonistic power, discriminatory preferences, rent-sharing or incentives heterogeneity across workers and estimation errors. The

[12]See Section 3 and Abowd et al. [1999] for a discussion of this regression model.

next paragraph discusses more explicitly the form of the labour supply function to the firm, providing methods that allow to control for the monopsonistic component of equation (3), *i.e.* the mark-down of wages with respect to the productivity induced by firms' factor market power.

2.3 The Labour Supply to the Firm

According to equation (3), in order to derive estimates of taste-based discrimination from firm-specific wage residuals, monopsonistic *mark-downs* of wages with respect to productivity have to be taken into account, at least unless we believe such mark-downs to be fairly close to zero. However, as mentioned in the Introduction, there are several reasons why employers might have some degree of wage-setting power in most labour markets. In particular, imperfect information about vacancies, direct and indirect costs associated to job switching, diseconomies of scale in hiring and employees monitoring costs, are some of the main reasons why the relevance of monopsonistic mechanisms should not be neglected.[13]

There are at least two special cases in which the task of correcting for monopsonistic mark-downs can be, from a computational perspective, quite simple. However, one limitation of these approaches is that for both cases discussed below it is necessary to identify a set of firms facing the same labour market structure, *i.e.* the same labour supply functions. The issue of how to construct groups of firms operating in homogeneous factor markets is discussed later in the paper. In particular, in Sections 4 and 5.2 we propose methods to select firms based on commuting zones, sectors of activity and observed job-to-job transitions. In this section, the availability of procedures allowing to identify *clusters* of firms facing a homogeneous labour market structure is taken as given.

As a first method to correct for mark-downs, assuming that firms face the same labour supply *function* (*i.e.* the same market structure), we consider an inverse labour supply

[13]See, among others, by Boal and Ransom [1997] and Manning [2003], for a detailed discussion of potential mechanisms providing employers with monopsonistic power against workers. Moreover, see Sulis [2011] for a direct assessment of the amount of labour market power held by firms in the Italian private sector, which is the market considered in the application of this model.

to the firm of the following form

$$w_i^g = (L_j^g)^{\alpha^g} (x_i^g)^{\beta^g} \tag{4}$$

where w_i^g is the individual wage expressed in levels, x_i is a vector of individual characteristics, a constant and an error term, β^g is a vector of parameters and two constant terms, L_j^g is days worked by gender g at firm j (where $i \in j$ is implicitly maintained) and α^g is a real-valued parameter. The characteristics included in the vector x_i^g control for all factors influencing individual productivity and individual heterogeneity in the supply of labour, so that the parameter α^g can be interpreted as a measure of elasticity of the labour supply specific of the given labour market, net of any other composition effect influencing the wage-size relationship.[14] If the above functional form is considered appropriate, this model can be written as a log-log regression of the form

$$\ln w_i^g = \alpha^g \ln L_j^g + \beta^g \ln x_i^g$$

The above function is a quite familiar wage equation, but provided that α^g is correctly estimated for each labour market, it can be given a structural interpretation. If this model is considered appropriate, correcting for mark-downs in equation (3) is straightforward, since for any group of firms facing the same labour market structure

$$\begin{aligned} \epsilon_j^g = \epsilon_s^g = \frac{1}{\alpha^g} \qquad & \forall\, s \neq j \\ \implies \ln\left(\frac{\epsilon_j^g}{1+\epsilon_j^g}\right) = \ln\left(\frac{1}{1+\alpha^g}\right) \approx -\alpha^g \qquad & \forall j \end{aligned}$$

It follows that if (4) is a correct functional form specification for the labour supply to the firm, then the gender wage gap in firm wage residuals can be corrected as

$$\widehat{\omega_j^m - \omega_j^f} = \omega_j^m - \omega_j^f + \alpha^m - \alpha^f \approx \hat{\delta}_j + \rho_j$$

[14] In the context of the previous section, $(x_i^g)^{\beta^g}$ is an element included in mp_i^g. This is consistent also with the empirical specification of the model, since all elements included in the vector x_i^g are also used to approximate for mp_i^g. In principle all factors *influencing only the demand of labour* (and not its supply) should be excluded from x_i^g, since such elements would bias the estimates of α^g towards zero.

In this case, since by assumption the gender-specific labour supply is a constant parameter across firms, correcting for mark-downs improves the precision in the estimates of the residual term $\hat{\delta}_j + \rho_j$, but it does not change the *relative ranking* of firms with respect to this component of the gender wage gap. However, as mentioned before this is generally true only if the gender-specific labour supply faced by firms is approximately constant across workplaces. That is, a labour structure of the form given in equation (4) provides an unbiased ranking of employers according to the intensity of their prejudices, as long as clusters of firms for which α^g is a constant parameter can be correctly identified.

In principle, there is at least another specification of the labour market structure that would be quite tractable in an empirical application. In particular, consider an inverse labour supply to the firm (again common across different workplaces) that reads as follows

$$w_i^g = \exp(\alpha^g L_j^g + \beta^g x_i^g) \quad \Longrightarrow \quad \ln w_i^g = \alpha^g L_j^g + \beta^g x_i^g \tag{5}$$

As before x_i contains individual characteristics accounting for heterogeneity in productivity and in the supply of labour, a constant and an error term. In this case, the labour supply elasticity to the firm can be written as

$$\epsilon_j^g = \frac{1}{\alpha^g L_j^g} \qquad \Longrightarrow \qquad \ln\left(\frac{\epsilon_j^g}{1+\epsilon_j^g}\right) \approx -\alpha^g L_j^g$$

It follows that whenever equation (5) provides a correct functional form representation for the labour supply to the firm, the gender wage gap in firm wage residuals should be corrected for monopsonistic mark-downs as follows

$$\widehat{\omega_j^m - \omega_j^f} = \omega_j^m - \omega_j^f - \alpha^m L_j^m + \alpha^f L_j^f \approx \hat{\delta}_j + \rho_j$$

In this case, correcting for mark-downs does not only affect the estimated magnitude of $\hat{\delta}_j + \rho_j$, but also the relative ranking of firms with respect to this variable. This is because now the wage elasticity faced by each firm depends (linearly) on the size of its male and female employment levels.

However, this last specification has some drawbacks. For example, the inverse linear relationship between firms' size and the elasticity of labour supply is imposed by construction. Moreover, the log-linear model is not very flexible if we want to try introducing more arbitrary heterogeneity in the elasticity of labour supply across clusters of firm's size, as this task would be quite demanding and prone to measurement error issues from a computational perspective.[15] Instead, under very similar assumptions to those imposed by the log-linear model, the log-log linear labour supply function allows to control for arbitrary differences in the elasticity across workplaces in a much simpler way, *i.e.* by adding fixed effects in equation (3).

The main advantage of a log-log functional form is indeed given by its tractability. In this case, it is more easy to control for differences in the labour supply elasticity across several dimensions, such as firms' employment levels, geographical areas and industries, with the only restriction for the elasticity to be approximately constant within so-defined clusters of similar firms. For this reason, in the empirical section of the paper we estimate a regression equation derived from the log-log linear inverse labour supply function, avoiding the use of linear, log-linear or linear-log specifications.[16]

Apart from considerations on what is the correct functional form specification for the labour supply to the firm, another difficulty in correcting for monopsonistic mark-downs is given by the presence of relevant confounding factors. For example, as documented by a vast stream of literature (*e.g.* Oi and Idson [1999]) there could be employer-size wage effects determined not only by monopsonistic mechanisms, but also by the fact that, for example, larger employers could attract workers of better quality, they could offer inferior working conditions, they could share a larger proportion of rents, or they may pay efficiency wages to deter shirking. Despite such limitation, employing the approach suggested by equation (4) is still useful in the present context. Indeed, the objective of our analysis is to derive a residual component of the gender wage gap corrected for monopsonistic mark-downs, which implies that the error induced by a spurious correlation

[15] For example, to introduce this kind of flexibility, it would be necessary to estimate pace-wise log-linear functions, such as that given by equation (5), on clusters of firms' size.

[16] Several of the considerations already mentioned for the log-linear case apply also to the linear or linear-log labour supply function model.

between firm size and wages (or between any of the categories used to partition firms and wages) is relevant for our purposes only as long as it differs by gender *and* across firms. This last point and other issues concerning the identification of the relevant parameters of the labour supply are further discussed in Section 5.2, where we present the empirical specification of the model in more detail.

2.4 The Role of Occupational Segregation

In the model of firms' wages given by equation (2), so-called *firms wage policies* contain two error terms, namely η_j, which is constant across workers of different sex within the same workplace, and ρ_j^g, which is gender- and firm-specific. One of the determinants of differences in the gender specific error terms (which we have previously denoted with ρ_j) is gender occupational segregation. The reason for this is that wage policies of employers could be different across type of workers. For example, the degree of rent sharing could be higher for those at the top of the occupational hierarchy. Similarly, the monitoring costs and working conditions could be different across occupations within the same firm, so that some employees could earn higher efficiency premiums or compensating wage differentials. If this was the case, any amount of gender occupational segregation could potentially contribute to the size of ρ_j.

In order to consider the possibility of heterogeneous rent-sharing (or similar mechanisms) across job titles, we assume for a moment that the workforce is divided into two occupational categories only (*e.g.* manual vs non-manual jobs), among which such heterogeneity is relevant. Since the focus of this section is on occupational segregation *within* firms, we omit the subscript j for the ease of notation. Let l_o^m and l_o^f be the employment levels in occupations $o = 1, 2$ for men and women. We begin by assuming that ρ^g is equal to zero, while η is different across occupations. More precisely, consider gender- and occupational-specific firm wage policies of the following form

$$\omega_o^g = \begin{cases} \ln\left(\frac{\epsilon^g}{1+\epsilon^g}\right) - \hat{\delta} 1\left[g = f\right] + c\eta & \text{if } o = 1 \\ \ln\left(\frac{\epsilon^g}{1+\epsilon^g}\right) - \hat{\delta} 1\left[g = f\right] + (1-c)\eta & \text{if } o = 2 \end{cases} \qquad 0 < c < 1$$

Given this definition, the average gender-specific firm wage policy can be written as

$$\omega^g = \ln\left(\frac{\epsilon^g}{1+\epsilon^g}\right) - \hat{\delta}1\left[g=f\right] + \left[\frac{l_1^g}{L^g}c + \frac{l_2^g}{L^g}(1-c)\right]\eta \qquad 0<c<1$$

Taking the difference between average pay premiums paid to men and to women, we have

$$\omega^m - \omega^f = \ln\left[\frac{(1+\epsilon^f)\epsilon^m}{(1+\epsilon^m)\epsilon^f}\right] + \hat{\delta} + \underbrace{\left[\left(\frac{l_1^m}{L^m} - \frac{l_1^f}{L^f}\right)c + \left(\frac{l_2^m}{L^m} - \frac{l_2^f}{L^f}\right)(1-c)\right]\eta}_{\equiv\,\rho}$$

Notice that the *segregation bias* ρ defined above can be expressed as a linear function of gender differences in the proportions of workers belonging to each occupation. Moreover, the term ρ tends to zero as the sum

$$\left|\frac{l_1^m}{L^m} - \frac{l_1^f}{L^f}\right| + \left|\frac{l_2^m}{L^m} - \frac{l_2^f}{L^f}\right|$$

tends to zero. The above sum can be interpreted as a *dissimilarity index*, which takes value zero if the relative proportions of men and women in each occupation are equal, while it takes value two if male and female workers are perfectly segregated into different occupations.[17]

The observations above can be used to generalize the discussion on the problem of occupational segregation. Suppose that employers share rents differently across a number O of distinct occupations. In this case, the comparison of male and female firm wage residuals provides a consistent representation of discriminatory preferences only under two circumstances. First, if differences in the proportion of workers belonging to each occupation are controlled for. Alternatively, this comparison is valid if the *occupational dissimilarity* index tends to zero. This index can be defined as

$$DI = \frac{1}{2}\sum_{o=1}^{O}\left|\frac{l_o^m}{L^m} - \frac{l_o^f}{L^f}\right| \qquad 0 \leq DI \leq 1$$

[17]See Graham [2017] for a recent discussion on the interpretation of the similarity index.

In the empirical analysis, we investigate how much several dimensions of gender segregation are relevant at the firm level. As discussed later in more detail, the possibility of heterogeneity in rent sharing across occupations is controlled for using detailed information on sector-specific job titles. Moreover, we test the robustness of our main results by excluding those firms where the gender dissimilarity index in occupational composition exceeds given thresholds.

3 Identification of Firm Wage Policies

In Section 2 we have shown that, given a constant labour supply across firms, a measure of taste-based discrimination can be recovered from gender differences in employer-specific wage residuals orthogonal to individual productive abilities. Given this theoretical implication, it is natural to employ a two-way fixed effects (or AKM) regression model (Abowd et al. [1999]) in order to identify firm wage premiums (or discounts) that employers grant to their workforce on top of pay components related to individual productivity.
On this respect, Card et al. [2016] have recently developed a useful approach, in the context of an AKM model, which allows to assess the *relative bargaining power* of women. More precisely, they show how much firm pay policies faced by men differ from women's ones, where such pay policies are defined as employer-specific wage components not related to characteristics equally rewarded across firms. Here, apart from some differences in the standardization of firm-specific wage residuals, we have followed quite closely their procedure.
Our identification strategy works as follows. Let i index a specific worker, t index the time, and $j = \iota(i,t)$ the firm in which i is working at t. Assume that the worker is observed for T time periods and let W_i represent a $T \times 1$ vector of daily wages, while X_i a $T \times P$ matrix of time- and firm-varying individual characteristics. Then, the two-way fixed effects model can be specified as follows

$$\ln w_{it} = x_{it}\beta + \eta_i + \omega_j + e_{it}$$

where w_{it} and x_{it} are rows of W_i and X_i respectively, β is a $P \times 1$ vector of parameters, while ω_j and η_i are respectively firm-constant and time-constant components of individual wages, which are allowed to be arbitrarily correlated with any of the characteristics in x_i, and which could be not perfectly observable.

The main assumption required for a consistent identification of the parameters is the absence of correlation between the error term e_{it} and all the other time-varying and time-invariant dependent variables. This condition must hold also for error terms in periods different from t, so that, for example, mobility towards employers with given firm wage policies can not be correlated with previous idiosyncratic shocks in earnings. Card et al. [2017] provide a careful discussion of this assumption, showing that overall the wage structure can be considered well approximated by the AKM additively linear specification.

The model is estimated separately among men and women. In both cases, the regressors included in the vector x_{it} are: a quadratic polynomial in age interacted with a dummy for part-time contracts; a quadratic polynomial in tenure interacted with a dummy for part-time contracts; a dummy for fixed-term contracts; four occupation dummies; year fixed effects. Moreover, all time-invariant and firm-invariant characteristics omitted from x_{it} are controlled for by the fixed effects included in the wage equation.

By computing the above regression model separately for men and women, we can obtain two estimators ($\hat{\omega}_j^m$ and $\hat{\omega}_j^f$) of the *gender-specific* firms' wage policies. Notice that these parameters are constant at the firm-gender level, and, as mentioned, they measure the additional wage premiums that some firms are willing to pay to their workers, independently of their characteristics. According to the model of Section 2, they represent a composite effect due to monopsonistic power, taste-based discrimination and other unobserved factors, such as rent-sharing or efficiency wages. The main step for isolating factors related to taste-based discrimination involves taking differences between male and female firm wage policies, recovering an expression for equation (3). The reminder of this section describes this procedure.

One of the main challenges in comparing $\hat{\omega}_j^m$ and $\hat{\omega}_j^f$ directly is given by the fact that both are computed with respect of an arbitrary reference group.[18] For this purpose, as discussed by Card et al. [2016], we need to define a firm, or a set of firms, as the common reference group across gender, and rescale each firm wage policy accordingly. The normalization choice proposed by these authors involves the use of balance sheet data, in order to set the lowest value-added group of firms as the reference. Here, we adopt a different choice, and select the largest firm (in terms of person-year observations) as the reference one. However, since the largest firm is different between female and male workers, to identify the common largest employer we define the following size function

$$size_j = \min\{\frac{N_{fj}}{N_{fj} + N_{mj}}, \frac{N_{mj}}{N_{fj} + N_{mj}}\}$$

where N_{jm} and N_{jf} are total firms' person-year observations among men and women, respectively. The reference firm that we have chosen is the largest according to the above definition of size, which gives more weight to workplaces in which the share of female and male workers is closer to 50%.

Denote by $\hat{\omega}^m_{\hat{j}}$ and $\hat{\omega}^f_{\hat{j}}$ the gender-specific pay policies of the reference firm $\hat{j}$ defined above. In order to normalize the firm pay policies, for any j we apply the following differences

$$\omega^m_j = \hat{\omega}^m_j - \hat{\omega}^m_{\hat{j}} \qquad \omega^f_j = \hat{\omega}^f_j - \hat{\omega}^f_{\hat{j}}$$

The normalization above allows to express all wage policies with respect to the ones applied by the largest one, for both men and women. Moreover, by considering the difference $\omega^m_j - \omega^f_j$, we can obtain an expression for equation (3) and test whether men' pay policy are proportionally higher (positive difference) or lower (negative difference) than women's pay policy at firm j, with respect to the same difference computed at firm $\hat{j}$. Finally, once that labour supply and segregation effects are controlled for, under reasonable assumptions the difference $\omega^m_j - \omega^f_j$ can be used to rank firms according to their *relative* discriminatory behaviour. Section 5.2 provides a discussion of this last feature of

[18] Also the number of reference groups depends on the number of *connected sets* among men and women (see Abowd et al. [2002]). For simplicity, we will restrict our analysis on the *largest connected set* only, which is composed of around 98% of all the observations.

the model.

4 Data and Sample Selection

The model discussed in Section 2 is estimated using Italian linked employer-employee data from administrative sources (Veneto Working History database, hereafter VWH). In particular, we study the population of private sector workers in the Veneto region of Italy during the period between 1996 and 2001. We have information on gross daily wages, inclusive of all pecuniary benefits paid by employers, some demographic and occupational characteristics, together with the location and the sector of activity of each firm.

In order to isolate the effects of taste-based discrimination, keeping constant any labour supply effects, we have to estimate the employer-specific gender wage gap, given by equation (3), on a very homogeneous sample of firms. More precisely, the main identifying assumption is that of a constant labour supply function across establishments. One way to approximate for this condition is to identify firms hiring from the same pool of workers. For this purpose, we have taken advantage of the comprehensive level of detail in the available data, exploiting also the peculiarities of the Italian region on which we focus our analysis.

Our strategy consists of three main steps. First, we identify *local labour markets*. Then, within such narrow regions, which we also call *districts*, we select manufacturing firms only. Finally, we select only those local labour markets where workers have particular mobility characteristics, in the sense that they tend to be employed only by firms within the district, and in which firms tend to hire from the same pool of workers, those belonging to their respective district.

Since 1991, the Italian statistical office (ISTAT) is in charge of identifying local labour markets (*Sistemi Locali del Lavoro*, or SLL). Districts definition is based on information gathered from the census and on a step-by-step procedure.[19] Using actual census data on individual commuting habits, a local labour market is defined as a group of highly

[19] See Lorenzini [2005] for the details of this procedure. A summary of the methodology employed in the identification of local labour markets from census data is available also online, see (`http://www.istat.it/it/files/2014/12/nota-metodologica_SLL2011_rev20150205.pdf`).

connected municipalities in terms of employment. Connectivity is defined using two main measures, the proportion of jobs within the districts held by its residents and the proportion of residents that work in the local labour market.

Using SLL as a first approximation, we then select only manufacturing firms.[20] This choice is motivated on two grounds. First, workers in the manufacturing sector tend to have much less geographical mobility. Similarly, firms in such sector tend to have a more homogeneous demand in terms of skills and they also tend to hire from the pool of workers available in a given labour market. Moreover, the Italian region under analysis, Veneto, is characterized by a large number of small and manufacturing-oriented firms, which tend to be located in particular areas of the region, forming high-density manufacturing conglomerates that tend to specialize in narrowly-defined activities. All of the above characteristics make manufacturing establishments an almost ideal unit of analysis, limiting the problem of confounding factors influencing the labour supply function faced by each firm.

In order to further limit problems related to firms' heterogeneity, we focus the analysis only on labour markets where: *i*) at least 60% of workers employed in the district never move to firms out of the SLL; *ii*) at least 30% of workers who change employer in the years of observation never move out of the district. To avoid using biased measures of firm wage policies, which are estimated as establishments' wage residuals in the context of an AKM regression model, we consider only firms where: *i*) at least 15% of total days worked in the firm are performed by one gender group; *ii*) 15% of workers employed by a firm over a 6-years period of observation are either men or women. Finally, since some of the local labour markets resulting from the above selection process are quite small, we restrict the attention on districts where at least 10,000 workers and 300 firms are observed over a 6-years period.

Table 1 summarizes the firms' selection criteria mentioned above and provides descriptive statistics on workers' mobility within the local labour markets included in the analysis. Two main statistics are reported, the ratio of workers changing district during the years of

[20] We have excluded from the analysis also two very marginal sectors, automotive production and oil refining. For these two-digit sectors, only around 20 firms where observed in the final sample.

Table 1: **Inclusion Criteria and Mobility Characteristics of Selected Districts**

Inclusion Criteria

Mobility	At least 60% of workers in the district are never observed employed at firms outside the district. At least 30% of workers changing firm never move outside of the district.	Gender composition	Only firms where at least 10% of workers over a 6-years period and 10% of days worked over a 6-years period belong to one gender group.
Sectors	Manufacturing sector only, with the exclusion of oil refining and automotive (*i.e.* very marginal sectors in the sample).	Size of Districts	At least 10,000 workers and 300 firms are observed in the district over a 6-years period (1996-2001).

Mobility Characteristics (Manufacturing-Sector Firms Only)

Included district name	**Number of Workers (1996-2001)**	**% Observed also Out of the District**	**% Changing Firm**	**% Observed also Out of the District Among Workers Changing Firm**
Padova	49,154	28.6	54.8	52.2
Verona	43,293	26.8	56.4	47.6
Vicenza	35,960	31.8	56	56.9
Treviso	35,086	34	55.9	60.7
Venezia	29,495	37.1	59.8	62.1
Arzignano	26,779	30.6	58.7	52.2
Conegliano	19,586	39.5	58.2	67.8
Montebelluna	17,205	39.7	57.5	69.1
Valdagno	15,930	28.2	42.4	66.6
Thiene	14,866	38	60	63.4
Bassano DG	14,431	35.4	57.7	61.4
Schio	11,865	34.1	54.9	62.1
Average Proportion	**Included Districts**	32.5		57.9
	Excluded Districts	35.6		63.4
P-val Equal Proportions		(0.000)		(0.000)

Table 2: **Descriptive Statistics by Gender in the Selected Sample**

	Women		**Men**	
	Mean	*St. Dev.*	*Mean*	*St. Dev.*
Log wage	4.609	0.319	4.883	0.369
Age	33.8	9.4	35.1	9.7
Tenure	6.6	7.1	6.4	6.9
Firm size*	7.2	29.2	10.6	39.5
Part-time	12.3%		1%	
Fixed-term	6.1%		4.6%	
Apprenticeship	5.1%		3.9%	
Blue collar	62.5%		69.9%	
White collar	32.4%		25.1%	
Manager	0.0%		1.1%	
N. firms	9,417		9,417	
N. firms-years	42,974		43,158	
N. workers	128,228		185,422	
N. workers-years	443,603		615408	

*: Firm size is computed as number of full-year equivalent workers by gender (total days worked in a year by a gender group within the firm, divided by 312). The average is taken considering a single observation per firm, without weighting for firms' size.

observation over the total number of employees in the SLL, and the proportion of workers changing district over the number of individuals observed at least once in the SLL and employed in more than one firm during the period of analysis. As can be noticed from the bottom part of the table, individuals in the included districts have a significantly lower tendency to move outside their respective local labour market than manufacturing workers in the excluded districts.[21]

Table 2 summarizes the main characteristics of the workforce by gender, considering only employees in the selected local labour markets. As can be noticed, the raw gender wage gap is of about 20%. Moreover, women are more likely to work part-time and in clerical occupations, they tend to be slightly younger and slightly over-represented among fixed-term contracts. On average, firms included in the analysis hire 7.2 women and 10.6 men, where these numbers refer to full-year contracts (*i.e.* 312 days). This last statistics reflects

[21]The same holds true also when comparing manufacturing and non-manufacturing firms, *i.e.* the mobility of workers in the selected sectors is less geographically dispersed. For brevity, this kind of analysis is omitted in the present context.

our inclusion procedure, through which only firms with a balanced gender composition of the workforce are considered in the analysis.

5 Empirical Results

In this section we present the main empirical findings of the paper. We begin by discussing the AKM regression results, through which gender-specific firm wage policies are estimated. Section 5.2 discusses the specification of a regression equation based on the gender wage gap in firms' pay policies. In particular, we show how the implied gender-specific labour supply to the firm can be taken into account and under which conditions a measure of taste-based discrimination net of monopsonistic considerations can be recovered.[22]

In Section 5.3 we present results on the effect of firm-specific gender pay differences implied by our regression model on traditional measures of the gender wage gap. We show that this parameter has a strong impact on differences in earnings between men and women. Moreover, its effects are linked to mechanisms that are not accounted for by differences in workforce composition nor by gender differences in returns to observable individual characteristics. Finally, Section 5.4 presents results on two proxies traditionally linked to taste-based discrimination, *i.e.* presence of a women at the top of a firm's hierarchy and share of female employment within firms. Our results provide strong evidences on the validity of both variables as proxies of discriminatory preferences. Moreover, we show that a second potential determinant of the firm-specific gender wage gap identified through our regression model is probably linked to within-firm differences in compensating wage differentials.

5.1 AKM Regression Results

We have estimated the AKM regression model separately by gender on the entire population of Veneto's private sector workers, considering the six-years period between 1996 and 2001. In order to estimate the model, we have selected one job spell per individual

[22] The Appendix A presents additional descriptive evidences on results based on this regression model.

Table 3: **AKM Regression Results by Gender in the Selected Sample**

	Women		Men	
$\mathrm{var}(\omega_j)$	0.022	22%	0.015	11%
$\mathrm{var}(x_{it}\beta + \eta_i)$	0.067	66%	0.103	76%
$2 * \mathrm{cov}(\omega_j, x_{it}\beta + \eta_i)$	-0.008	-7%	0.009	6%
$\mathrm{var}(e_{it})$	0.020	20%	0.008	6%
$\mathrm{var}(w_{it})$	0.102	100%	0.136	100%
N. Observations	443,603		615,408	

The table presents the wage variance decomposition based on the AKM regression model. The parameters of the regression are estimated separately by gender on the entire database of Veneto's private sector. The results of the table are instead computed considering only the sample of manufacturing firms selected along the lines discussed in Section 4. The percentages reported in the table are expressed in terms of the total gender-specific wage variance.

in each year, choosing the longest work episode whenever a person was employed at more than one firm in a given year. Finally, we have restricted the analysis on firms belonging to the largest connected set, *i.e.* the set of all establishments connected by the mobility of workers.[23]

The reason for estimating the regression model on the entire sample of Vento's private sector workers is given by the fact that firm wage policies are measured with more precision the largest the number of firm-to-firm mobility episodes observed in the data. Having estimated the AKM regression for all firms of Veneto, we then analyse its parameters considering only manufacturing establishments belonging to one of the twelve local labour markets satisfying the inclusion restrictions discussed in Section 4.

Table 3 summarizes the results of the regression model, separately by gender, on the sample of manufacturing firms belonging to one of the selected districts. It can be noticed that overall wage dispersion is higher among men. In both cases, the largest contribution to total wage dispersion is given by the joint effect of individual time-varying and time invariant characteristics and by the returns to such endowments. Moreover, the regression residual is larger among women, implying that the model fits better the data in the case of men.

[23] This set corresponds to around 98% of the observations. See [Abowd et al., 2002] for a discussion of this procedure and a more detailed definition of connected sets.

Figure 2: **Comparison of Men's and Women's Firm Wage Policies by Sector**

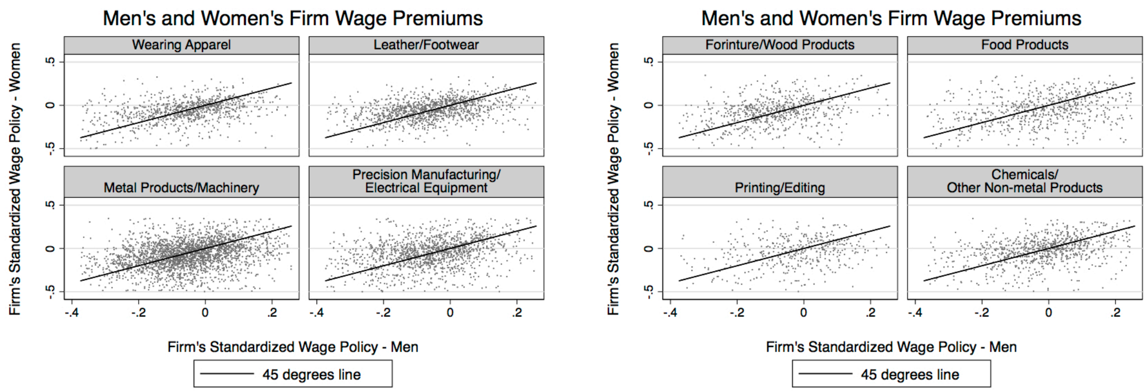

Firms' wage policies are expressed as a difference from those paid by the largest employer of both gender groups. For firms on the 45 degrees line the difference between wage policies paid to men and women is the same as that observed at the reference firm. For graphical convenience, outliers above or below 1.96 standard deviations from the gender-specific average of firms' pay policies are omitted.

In the context of the present analysis, the two most interesting elements of earnings variability are represented by the variance of firms' pay policies and by the sorting of these residuals with individual *portable* wage components. Firm wage premiums provide a larger contribution to wage dispersion for women (22%) than for men (11%). In general, this result is consistent with our theoretical model, given that taste-based discrimination represents an element of variability in firm wage policies that is absent in the case of men. However, this result could also be partly driven by a larger measurement error in firm-specific wage residuals. Moreover, it is interesting to notice that while the sorting of highly paid workers to highly paying firms is positive in the case of men, it is instead negative in the case of women. Also this result could be considered consistent with our theoretical model, given that firm's wage policies paid to women are distorted by taste-based discrimination. However, as suggested, among others, by Andrews et al. [2008], one should be cautious in interpreting this correlation, given that measurement error in either firm wage policies or individual productivity is likely to induce a negative bias to this parameter.

We next standardize firm wage policies paid to men and women in order to make them comparable. As discussed in Section 3, this task is performed by choosing a common reference firm, which we decide to be the largest employer of both gender groups (*i.e.* the one employing the largest number of workers belonging to the gender group less represented within the firm).

Figure 2 compares firms' standardized wage premiums paid to men with those paid to women. Establishments are classified into eight manufacturing sectors and each dot represents one of them. Firms lying below the 45 degrees line pay relatively higher premiums to men than to women with respect to what is observed at the reference firm. The opposite holds true for observations above the 45 degrees line. In this sense, the graphs provide a *relative* measure of the gender wage gap in firms' residuals, not an absolute one.[24]

A general pattern that emerges from the figure is that as men's pay policies (on the horizontal axis) grow, the density of firms giving a relatively better treatment to men also grows. Instead, for firms paying relatively lower pay policies to men, the density of firms above the 45 degrees line tends to grow. This general pattern is consistent with the findings of [Card et al., 2016], who show on Portuguese data that as firms' rents (and consequently also wage policies) grow, the gender wage gap in pay premiums tends to widen. However, some degree of heterogeneity across sectors also emerges. For example, the density of firms well below the 45 degrees line seems to be relatively higher in metal-manufacturing and machinery production sectors than, for example, in leather and footwear industries.[25]

Figure 2 provides a similar comparison of firm wage premiums between men and women, this time displayed by local labour market. Also in this case it emerges a general flat pattern, in which the gender wage gap grows larger the higher men's pay policies. However, the cloud of observations has some degree of heterogeneity across district, as it tends to be relatively more dense around or above the 45 degrees line in some cases (*e.g.* Conegliano, Valdagno, Montebelluna) or relatively more spread toward the bottom-right of the panel in other cases (*e.g.* Arzignano, Venezia, Thiene, Schio), even if, by simple

[24] Notice that observations above or below the 45 degrees line should not be interpreted as having a *negative* or *positive* gender wage gap, as this measure can only be expressed with respect to a reference firm or group of firms. See Card et al. [2016] for a detailed discussion of this point.

[25] The average gender wage gap in standardized firms' residuals is significantly lower (at a 5% significance level) than in the rest of the sample for the following sectors: leather and footwear, furniture/wood products, precision manufacturing/electrical equipment. It is significantly higher in the following industries: metal products/machinery, food products.

Figure 3: **Comparison of Men's and Women's Firm Wage Policies by District**

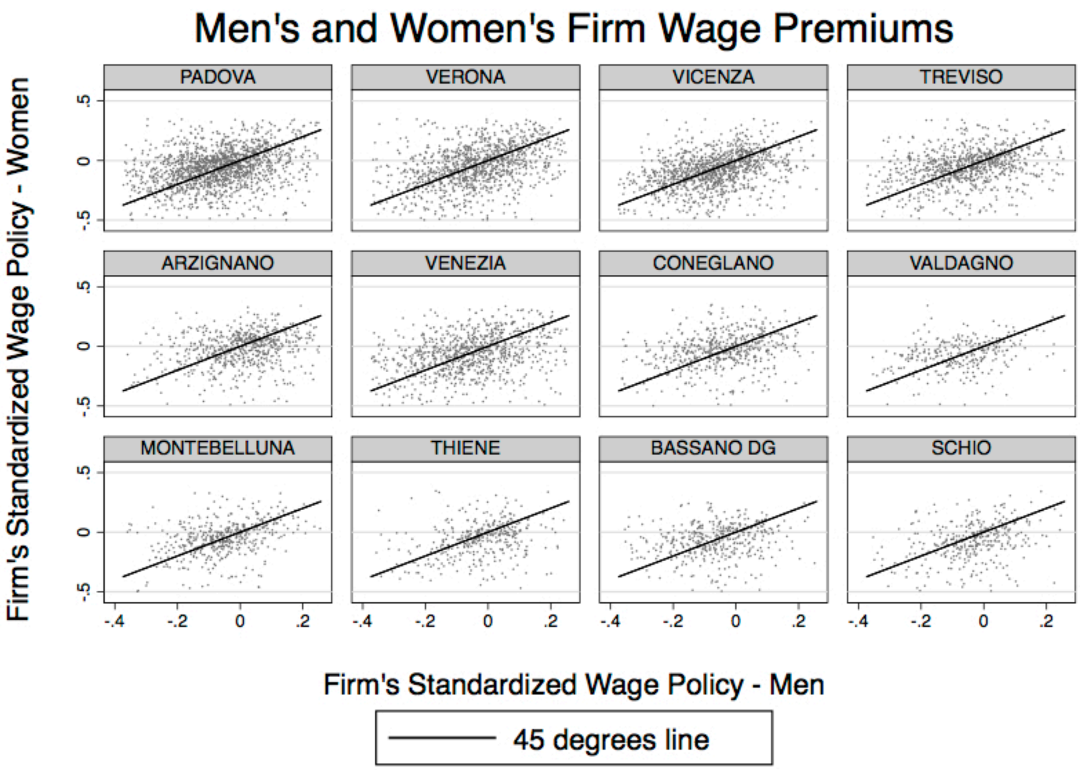

Firms' wage policies are expressed as a difference from those paid by the largest employer of both gender groups. For firms on the 45 degrees line the difference between wage policies paid to men and women is the same as that observed at the reference firm. For graphical convenience, outliers above or below 1.96 standard deviations from the gender-specific average of firms' pay policies are omitted.

eyeball inspection, overall differences are not always neat.[26]

5.2 Specification of the Firm-Specific Gender Wage Gap Equation

We now turn to the problem of identifying employers' discriminatory tastes starting from the gender wage gap in firms' wage policies. The presence of several possible confounding factors should be acknowledged in this exercise, and we begin by discussing the role of monopsonistic discrimination. As shown in Section 2.3, from a theoretical perspective firms' premiums can be influenced by the degree of factor market power held by employers. Since mark-downs of wages with respect to productivity could differ by gender and across firms, employers' discriminatory preferences against women can be recovered only by controlling for heterogeneities in firms' pay policies induced by monopsonistic mechanisms, *i.e.* by differences in the labour supply to the firm.

Following the discussion of Section 2.3, we assume that each firm belongs to a given labour market k. Within such markets, all employers face an identical gender-specific log-log linear inverse labour supply function that read as

$$\ln w_i^g = \alpha_k^g \ln L_j^g + \beta_k^g x_i^g \qquad g = m, f \quad k = 1, \ldots, K$$

where the vector x_i^g contains all characteristics affecting individual labour supply and an error term. Given this functional form, it follows that the labour supply elasticity to the firm is determined by α_k^g only. Moreover, the above equation implies that the gender wage gap in firms' wage residuals can be modelled as follows

$$\omega_j^m - \omega_j^f \approx \alpha_k^f - \alpha_k^m + \hat{\delta}_j + \rho_j^m - \rho_j^f \tag{6}$$

where $\hat{\delta}_j$ represents discriminatory preferences against women, while ρ_j^g is a firm- and gender-specific residual attributable to heterogeneous rent-sharing, compensating wage differentials, incentives or measurement error between men and women. Before present-

[26]The average gender wage gap in standardized firms' residuals is significantly lower (at a 5% significance level) than in the rest of the sample for the following districts: Conegliano, Padova. It is significantly higher in the following districts: Arzignano, Schio, Thiene.

ing results on regression models implied by equation (6), something on which we dedicate the next two paragraphs, we first discuss the main identifying assumptions of our approach and provide more details on our specification choices.

As already mentioned, one of the main assumptions we need to make is that, within a given market k, the gender-specific inverse labour supply function has a constant derivative across firms and it is well approximated by a log-log functional form. If this condition holds, then the inclusion of fixed effects for each labour market k allows to control for the term $\alpha_k^f - \alpha_k^m$ in equation (6). A related assumption of the model is that discriminatory tastes are not correlated across labour markets. If this latter assumption is violated, then interpreting differences in the gender wage gap across markets becomes more problematic, given that fixed effects included in the regression equation would control also for genuine differences induced by higher or lower taste-based discrimination. Nevertheless, a violation of this last assumption would not affect the internal validity of our results for a given labour market.

Whether clusters of firms facing the same labour market structure can be identified correctly remains a matter of judgement. As a first approximation, we consider only the geographical districts of Veneto selected along the lines discussed in Section 4. However, there are still reasons to believe that relevant heterogeneities in the labour supply elasticity to the firm persist even within such local labour markets. For example, firms operating in different industries could be hiring from different pools of workers. Moreover, the same could be true also for relatively large or small firms, *e.g.* bigger establishments could be able to affect their labour supply and/or they could demand different sets of skills and qualifications. For this reason, we include in the regression equation a finer set of fixed effects than just those controlling for local labour markets.

In particular, we include in equation (6) a full set of dummies for each two-digits sector. Furthermore, we include in the regression model also a full set of dummies controlling for each quartile of the firm size distribution, where establishment size is approximated by average total days worked within the firm in a year and percentiles are defined over the distribution of firms without weighting for the number of employees. For robustness, we test our main results also using a more saturated model in which firm size and sector are interacted with dummies controlling for each local labour market.[27] Moreover, we test whether results differ when including also fixed effects at the municipal level together with those at the local labour market level.[28]

A second source of potential bias in identifying the parameters of equation (6) arises from the error term $\rho_j^m - \rho_j^f$. In principle, several confounding factors could weaken our results whenever discriminatory tastes are correlated with elements included in such gender-specific firms' wage residuals. One particularly important case, that we have discussed in Section 2.4, is given by the presence of heterogeneous rent-sharing across occupations. In order to control for this source of bias, we have used information on a detailed list of job titles defined and protected by collective bargaining institutions.

Approximately once every two years Italian trade unions and employers' associations agree on a set of rules establishing a classification of workers according to their occupation. Such mandatory rules define workers' tasks and qualifications in each of these job titles, setting also occupation-specific minimum wage standards. Thus, such job titles (called *livelli di inquadramento* in Italian) represent an ideal categorization of the workforce, which allows to estimate the extent of occupational segregation in great detail. However, there are also some drawbacks in using *livelli di inquadramento*. In particular, the definition of job titles applies at the industry-wide level and is not harmonized across sectors. Moreover, it may change from year to year. For these reasons, and given that a substantial amount of errors in reporting the correct classification code were detected in the data,[29] we have aggregated all *livelli di inquadramento* into five percentile groups for each two-digit sector, where such percentiles are defined on the year- and sector-specific distribution of average

[27] However, a more saturated model in which firm size is interacted with local labour markets does not provide a much better fit to the regression equation, given that the relationship between firm size and the gender wage gap is quite similar across districts. The same holds true also for district-industry interactions.

[28] Municipalities are the smallest administrative entities on the Italian territory. The local labour markets included in our analysis are composed of a number of municipalities that varies between 6 (in Valdagno) and 49 (in Padova).

[29] Job titles had to be imputed for around one-third of the observations, which did not match with a database containing the exact classification code of occupations. The adopted imputation algorithm assigns to each missing observation the (non-missing) job title of the worker earning the closest daily wage on the two-digits sector- and year-specific pay distribution.

wages observed within each job title.

Considering the above mentioned sector-specific occupations, we have included as controls in the regression equation gender differences in the proportion of workers belonging to each job titles within the firm. In particular, we have defined this variable to be positive whenever the proportion of men in a given job title within the firm was higher than the same proportion among women. Moreover, such variable takes negative values in the opposite case and it is equal to zero if no worker is observed in a given job title within the firm. For robustness, we have also tested the main results by excluding firms in which the occupational dissimilarity index (defined in Section 2.4) exceeds given thresholds.[30]

The Appendix A provides several descriptive statistics on the residual term $\hat{\delta}_j + \rho_j^m - \rho_j^f$, together with a discussion on how the inclusion of covariates that approximate for $\alpha_k^m - \alpha_k^f$ and for $\rho_j^m - \rho_j^f$ affects its size and distribution. Instead, in the next two paragraphs we discuss results derived from the regression model implied by equation (6) following different approaches. In Section 5.3 we measure the impact of the residual component $\hat{\delta}_j + \rho_j^m - \rho_j^f$ on traditional measures of the gender wage gap. Section 5.4 proposes a parametrization of this residual, in which proxies traditionally associated to taste-based discrimination enter in $\hat{\delta}_j$ and can be tested empirically.

5.3 Overall Impact of the Firm-Specific Gender Wage Gap

In this section, we evaluate to which extent traditional measures of the gender wage gap can be explained by the residual term of equation (6). Results presented here may be interpreted as a "potential impact" of taste-based discrimination, but this approach has also obvious limitations. As suggested by its definition, the residual term $\hat{\delta}_j + \rho_j^m - \rho_j^f$ may be affected by several confounding factors. Therefore, the estimates presented in this section could represent only a lower or an upper bound of the true effects of discriminatory preferences, depending on whether such tastes are positively or negatively correlated with firm-specific gender differences in rent-sharing, compensating wage differentials or measurement errors. Despite such limitation, the analysis of this paragraph can be never-

[30] As discussed in Section 2.4, whenever the gender dissimilarity index in occupations tends to zero, so does the bias induced by gender segregation within the firm.

theless highly informative, since it sheds light on the quantitative importance of a residual term on which the existing literature, apart from the seminal contribution of Card et al. [2016], has seldom focused. The findings presented here also provide a useful framework for interpreting the results of the next section, in which we carry on the more rigorous exercise of testing a parametrization of $\hat{\delta}_j$.

In order to quantify the effect of being employed at a more discriminatory firm on the gender wage gap, we define the following treatment variables

$$T_\theta = 1\left[g = m\right] * 1\left[\mathrm{F}(\hat{\delta}_j + \rho_j^m - \rho_j^f) > \theta\right] \qquad \theta = \{10, 25, 50, 75, 90\}$$

where $1\left[g = m\right]$ is a dummy for male workers and F() is the cumulative distribution function (over firms) of the variable $\hat{\delta}_j + \rho_j^m - \rho_j^f$. This last term is defined as the residual of an OLS regression of $\omega_j^m - \omega_j^f$ on: district fixed effects, two-digits sector fixed effects, class of firm size fixed effects and gender differences in the proportion of job titles within the firm.

By estimating the effect of the treatment T_θ in a cross-sectional log wage equation, we have a measure of how much the gender wage gap grows among employees at firms in the right tail of the "discrimination distribution", keeping fixed observable individual characteristics. In particular, for each quantile θ of F() we estimate the following model

$$\ln w_i = b_1 1\left[g = m\right] + b_2 T_\theta + \beta x_i + \eta_j$$

where η_j is a firm fixed effect (common for both gender groups), while x_i is a vector of controls, a constant and an error term. We estimate the above model by OLS on the cross-section of workers in 2000, which is the year with most observations in our sample, clustering standard errors at the establishment level. We test two sets of independent variables. The baseline model (Model 1) includes: firm fixed effects, quadratic polynomials in age and tenure, a dummy for part-time interacted with both of these polynomials, four main controls for occupation (*i.e.* apprenticeships, manual workers, clerical workers and managers) and a full set of job titles fixed effects, where such job titles are sector-specific and defined from collective bargaining classifications along the lines discussed in the previous section.[31] The second model that we test (Model 2) interacts all independent variables previously included in x_i with a dummy for gender.[32]

Given these specifications, the treatment effect in Model (1) is interpreted as a growth in the part of the gender wage gap that is not explained by individual characteristics included in the regression. Instead, Model (2) allows the returns to endowments to be gender-specific, so that b_2 in this case is interpreted as an additional effect of T_θ on the gender wage gap controlling for both, *characteristics effects* and *coefficient effects* typically arising in Oaxaca-Blinder decompositions (see Oaxaca [1973]). Notice also that since firm fixed effects are included in both models, the treatment effect is net of any wage premium associated to being employed at a given establishment, as long as this premium is equally affecting men and women.

Figure 4 presents the results obtained by estimating the models discussed in this section. The graph in the top panel shows, considering both the baseline and the fully interacted specifications, how the coefficient b_2 varies when T_θ is defined using different percentiles of the distribution F(). In general, the treatment effect is always strong and significant. Women employed at more discriminatory firms suffer an additional wage loss with respect to men of between 4% and more than 10%, depending on how the treatment variable is defined. This implies that the gender wage gap conditional on individual characteristics grows substantially with respect to its baseline level. For example, in the baseline model this growth is of more than 40% when θ is equal to 0.5.

Treatment effects are not constant as the definition of more discriminatory establishments changes. In particular, as the definition of this group shifts toward the right tail of the distribution F(), the wage penalty faced by women at such establishments grows.[33] It is interesting to notice that treatment effects are almost the same when comparing the fully saturated model with the baseline specification. This result is nevertheless less surprising

[31] In particular, we control for five job title fixed effects for each two-digit sector, which amounts to a total of 95 job titles fixed effects.

[32] Obviously, an interaction of firm fixed effects with a gender dummy was not included, as it would have been collinear with the treatment variable T_θ.

[33] This result should however be interpreted with caution, as outliers in the right tail of F() are given more weight as θ grows. The Appendix A presents several plots of the distribution of the term $\hat{\delta}_j + \rho_j^m - \rho_j^f$.

Figure 4: **Impact of the Residual GWG in Firms' Premiums on the Cross-Sectional GWG (Year 2000)**

Marginal Effect of Being in the Right Tail of Most Discriminatory Firms

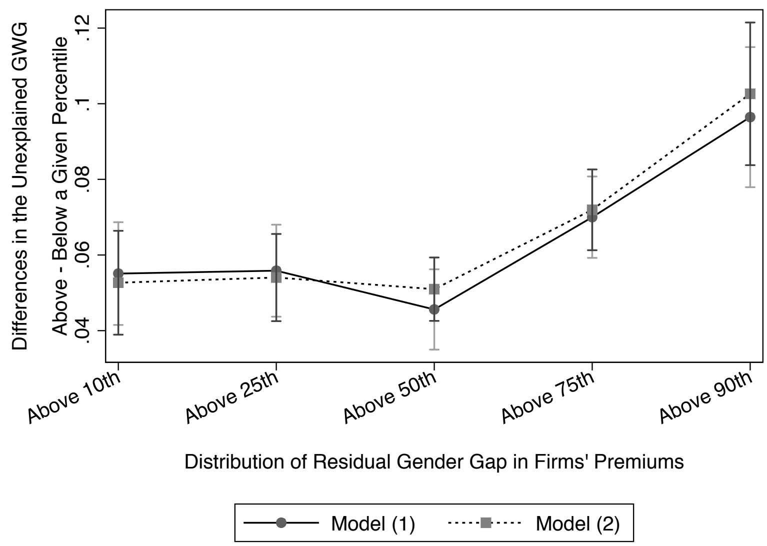

Summary of Regression Results

Effect of Being Above the 50th Percentile of Most Discriminatory Firms

	Dependent variable: log daily wage	
	Model (1)	**Model (2)**
Coefficients		
$1\,[g=m] * 1\left[\mathrm{F}(\hat{\delta}_j + \rho_j) > 0.5\right]$	0.046^{**}	0.051^{**}
95% CI	$[0.035, 0.056]$	$[0.043, 0.059]$
$1\,[g=m]$	0.111^{**}	-0.033
95% CI	$[0.104, 0.118]$	$[-0.205, 0.140]$
F tests		
Age and tenure polyn.	233.37^{**}	213.76^{**}
Interactions with $1\,[g=f]$		67.36^{**}
Part-time and interactions	578.01^{**}	159.84^{**}
Interactions with $1\,[g=f]$		4.42^{**}
Main occupation dummies	718.43^{**}	525.82^{**}
Interactions with $1\,[g=f]$		17.23^{**}
All covariates	723.20^{**}	557.40^{**}
Adjusted R^2	0.751	0.755
RMSE	0.193	0.191
N. firm effects	7,564	7,564
N. of observations	178,964	178,964

S.e. clustered by firm. Significance levels: $**$*: 1%;* $*$*: 5%*

Results of regressions of log wages in year 2000 on a dummy for male workers at firms in the right tail of the distribution F() of $(\hat{\delta}_j + \rho_j)$. The graph in the top panel shows treatment effects using different percentiles of F(). The table summarizes results of the regression models when the treatment is being male and above the median of F(). Model (1) includes standard controls, Model (2) interacts all dependent variables (apart from firm fixed effects) with a gender dummy.

once that we consider the definition of $\hat{\delta}_j + \rho_j^m - \rho_j^f$. Indeed, this residual captures gender differences in wages net of individual observable and time-constant characteristics and firm effects that are common for both gender groups. Thus, this variable represents an effect that is usually left unexplained in traditional cross-sectional estimates of the gender wage gap. As can be noticed, even allowing returns to endowments and age earning profiles to be gender-specific does affect the significance and magnitude of these treatment effects.

In the table below the graph, we provide more detailed information on the regression model, considering only the case in which θ is set to 0.5. As can be noticed, the estimation sample is composed of almost 180,000 workers and around 7,500 firm fixed effects are included in the regression. The F-tests show that the main independent variables are jointly significant and the overall goodness of fit of the regression is quite high in both models. In the baseline specification, the gender wage gap grows from around 11% to almost 16% for female workers in the top 50% of most discriminatory firms. When returns to observable individual characteristics are allowed to be gender specific, the coefficient b_1 is no more significant, as its effect is fully absorbed by such gender-specific coefficients. However, the same does not hold true for the treatment effect b_2, as this coefficient is almost equal in the two models.

Results presented in Figure 4 are interesting, as they underline the fact that the firm-specific residual gender pay gap is associated with elements that in a traditional cross-sectional analysis can only be ascribed to unexplained *coefficient effects*. However, as mentioned before these results are also difficult to interpret, given the ambiguity on which component of the term $\hat{\delta}_j + \rho_j^m - \rho_j^f$ is the most relevant and which mechanisms are producing the most important effects on the gender wage gap. In the next section we discuss a more rigorous approach that, by parameterizing this residual, aims at overcoming this limitation of our analysis.

5.4 Parameterization and Estimation of Taste-Based Discrimination

In this section we test a parametrization for $\hat{\delta}_j$, through which we aim at assessing whether proxies traditionally associated to taste-based discrimination have a theoretically coherent impact on the firm-specific gender wage gap, and can thus be considered valid. In particular, we test the following two proxies: presence of women at the top of the firm's hierarchy and female share of workers within firms.

As shown in the theoretical model, the fact that more discriminatory employers hire a lower share of female workers holds always true only if *labour supply functions and marginal revenue product curves are kept fixed across firms.* Whenever these two comparative statics conditions do not hold, in principle it could be possible for employers with higher taste-based discrimination to also have a higher share of female workers within the firm.[34] Nevertheless, given that discriminatory firms hire less women than what would be observed at their monopsonistic benchmark, it may be reasonable to expect that this variable is at least correlated with preferences against women.

The second proxy that is included in $\hat{\delta}_j$, *i.e.* presence of women at the top of the firm's hierarchy, is based on the assumption that if women are represented in managerial positions, then they are more able to affect the firm's culture and to prevent discriminatory behaviours. Several contributions have investigated the implications of Beker's or similar discrimination theories using one of these two proxies.[35] However, to the best of our knowledge, a formal tests on whether it is correct to take these variables as firm-specific measures of taste-based discrimination has never been constructed.

In order to control for potential confounding factors, we test further mechanisms that could enter in $\rho_j^m - \rho_j^f$ and be correlated with our proxies of taste-based discrimination, focusing in particular on compensating wage differentials. From a theoretical perspective, compensating wage differentials may allow firms with given characteristics to hire women at a wage below the equilibrium level. Here we test the assumption that firms that offer a higher proportion of part-time contracts are able to hire women paying them less, due to the fact that female employees prefer more flexible working schedules. Thus, we use

[34] In Figure 1, the female share depends not only by δ, but also by the horizontal distance between the male and female marginal factor cost curves.

[35] For example, Flabbi et al. [2014] study the impact of the presence of women in managerial positions on firm performance and the gender wage gap, assuming that female managers have different attitudes toward female workers. Similarly, Weber and Zulehner [2014] study the impact of the employment share of women within firms on their survival probability, assuming that such workplaces, being less discriminatory, are able to better survive competition due to their more efficient decision-making.

weeks worked part-time over total weeks worked within the firm in a year as a proxy for this mechanism.

Starting from equation (6), we consider the following regression model

$$\omega_j^m - \omega_j^f = \alpha_k + \underbrace{b_1\hat{\delta}_j^1 + b_2\hat{\delta}_j^2 + \hat{\delta}_j^r}_{\equiv\hat{\delta}_j} + \underbrace{b_3\rho_j^1 + \rho_j^r}_{\equiv\rho_j^m - \rho_j^f}$$

where, in the baseline specification, α_k is a vector of fixed effects for district, two-digit sector and class of firm size. ρ_j^r is a vector containing five variables on gender differences in the proportions of job titles within the firm and the ratio of weeks worked part-time over total weeks worked within the firm in a year.[36] $\hat{\delta}_j^r$ is a term for the residual amount of taste-based discrimination in a given firm and $\hat{\rho}_j^r$ is a residual error term, where the sum of these residuals has mean zero and is assumed to be not correlated with all the other variables in the model. The two independent variables of most interest are $\hat{\delta}_j^1$, which is equal to one if there is a women in the management of the firm, and $\hat{\delta}_j^2$, representing the ratio of total days worked by women over total days worked within the firm during the entire period of observation.

Since information on firms' ownership and management structure was not available in the data, we have defined as female managers all women in non-manual occupations that were receiving the highest observed yearly earning within the firm, where the highest pay is defined over all person-year observations in the period 1996-2001. For firms with more than 60 person-year observations, we have relaxed this definition and considered as female managers also those women in non-manual occupations that were among the top 3% yearly income earners *and* one of the top 10 earners among all person-years observations of a given workplace.

A difficulty in measuring the effects of $\hat{\delta}_j^1$ and $\hat{\delta}_j^2$ on the firms' premium gender gap is that both variables may be correlated with the error term. In particular, in the case of female managers the problem may be one of simultaneity, given that firms' wage policies

[36]This variable is computed year-by-year within each workplace. It is then averaged over the years in which a firm is observed

Table 4: **Summary Statistics on Proxy Variables**

Variable	**Mean**	**St. Dev.**	**N. Obs.**
Female manager	13.3%		9,417
Female manager (1992-1995)	14.4%		7,855
Female share	0.439	0.198	9,417
Female share (1992-1995)	0.418	0.247	7,855
Part-time share	0.056	0.089	9,417

All statistics are computed over firms. The number of observations for variables measured in the period 1992-1995 refers to all firms that could be merged with the 1996-2001 sample.

are probably higher for women in workplaces where one of them is earning the most.[37] Similarly, the ratio of female over total employment could be a variable correlated with measurement errors in monopsonistic discrimination or in firms' wage policies themselves. For these reasons, taking advantage of the long panel structure of the VWH database, we test the validity of our results by measuring these two variables also over the period 1992-1995, looking at whether their relationship with $\omega_j^m - \omega_j^f$, as measured over the period 1996-2001, still holds. The underlying assumption made in this exercise is that a firms' culture and its management structure change only slowly over time. If this is the case, then the advantage of measuring the variables $\hat{\delta}_j^1$ and $\hat{\delta}_j^2$ during the period 1992-1995 is that they can still be good proxies for attitudes toward women also in the subsequent years, while at the same time they are less vulnerable to the problem of simultaneity and correlation with the measurement error.

Table 4 shows descriptive statistics on $\hat{\delta}_j^1$ and $\hat{\delta}_j^2$, as measured both, over the period 1996-2001 and the years 1992-1995, together with a summary of the proxy for compensating wage differentials. As can be noticed, only around 83.4% of the firms included in our sample could be observed also during the years 1992-1995.[38] The proportion of firms managed

[37]Notice however that firms' pay premiums are estimated conditioning on human capital characteristics of individual workers. Thus, in our context the problem of simultaneity is attenuated by this feature of the AKM regression model.

[38]Given that the VWH database covers the *population* of Veneto's private sector firms, this attrition can be ascribed only to either the opening of new plants or to spurious discontinuities in reporting the firm identifier from year to year. However we consider this last mechanism arguably quite rare, given that firm identifiers have an important purpose from an administrative point of view.

by a women is of 13.3% in 1996-2001, while the same number is slightly higher (14.4%) when measured during previous years. The proportion of days worked by women within the firm is also quite high, something that partly reflects the sample selection choices that we have adopted for the analysis. The same variable is indeed smaller for firms that could be observed over the period 1992-1995. Finally, the proportion of weeks worked under part-time contracts is quite low in our sample.

The regression results are presented in Table 5. All the parameters are estimated by OLS clustering standard errors at the municipality and industry level. In the baseline specification (Model 1), all variables are measured over the period 1996-2001 and only the above mentioned main controls for occupational segregation, size, district and industry are included. As can be noticed, the presence of a female manager within the firm negatively affects the gender wage gap in firms' premiums. The same holds true also for what concerns the female share of workers, while firms offering more flexible work arrangements, coherently with what a model of compensating wage differentials would suggest, tend to pay women less. Even though the F-statistics show that the main controls of the model are jointly significant, the overall regression fit is quite low, something that is nevertheless not surprising when considering the fact the dependent variable is a difference between two regression residuals.

The second column of results (Model 2) refers to a regression in which all controls are kept equal to Model 1, except for $\hat{\delta}_j^1$ and $\hat{\delta}_j^2$ that are now measured over the years 1992-1995. Interestingly, the effects of both these variables are still negative and significant, suggesting that they are good proxies of taste-based discrimination. Moreover, this is arguably a quite solid result, which is not simply driven by simultaneity or measurement error issues. In Model 3 and 4 we further test the robustness of our results, by increasing the flexibility of the regression model and by allowing a larger number of unobservable effects to be held fixed.

In Model 3 we interact gender differences in the proportion of job titles with two digits sector fixed effects, in order to better control for within-firms gender segregation. Moreover, we interact dummies for class of firms' size with local labour market fixed effects,

Table 5: **Regression Results on Taste-Based Discrimination**

	Dependent variable: GWG in firms' premiums			
Model	**(1)**	**(2)**	**(3)**	**(4)**
Coefficients				
Female manager	-0.027^{**}			
95% CI	[-0.040,-0.013]			
Female manager (1992-1995)		-0.014^{*}	-0.014^{*}	-0.021^{**}
95% CI		[-0.027,-0.001]	[-0.028,-0.001]	[-0.036,-0.006]
Female share	-0.065^{**}			
95% CI	[-0.094,-0.037]			
Female share (1992-1995)		-0.062^{**}	-0.06^{**}	-0.055^{**}
95% CI		[-0.086,-0.038]	[-0.085,-0.034]	[-0.082,-0.028]
Part-time share	0.126^{**}	0.129^{**}	0.142^{**}	0.169^{**}
95% CI	[0.050, 0.202]	[0.047, 0.21]	[0.058, 0.225]	[0.071, 0.267]
F tests				
Firm size f.e.	13.51^{**}	8.79^{**}		
Firm size∗district f.e.			1.76^{*}	1.82^{**}
Segregation	3.36^{**}	3.62^{**}		
Segregation∗sector f.e.			1.74^{**}	1.82^{**}
Sector f.e.	2.73^{**}	2.36^{**}		
District f.e.	2.16^{*}	2.81^{**}		
District f.e.∗sector f.e.			2.78^{**}	3.07^{**}
Municipality f.e.			2.88^{**}	3^{**}
All covariates	5.31^{**}	4.84^{**}	22.19^{**}	30.17^{**}
Adjusted R^2	0.017	0.020	0.028	0.025
RMSE	0.225	0.213	0.212	0.198
N. of observations	9,417	7,855	7,855	6,084
Clusters for s.e.	Two-digits sect.∗municip.		Three-digits sect.∗municip.	

Significance levels: ∗∗: 1%; ∗: 5%

The table summarizes the results of regressions of proxies for taste-based discrimination and compensating wage differentials on the gender wage gap in firms' pay policies. The units of observation are individual firms. Basic controls include measures of gender occupational segregation, industry, geographical area and size (yearly average for each firm, controlled for using four dummies for each quartile of its distribution in the regression). Model 2 uses proxies for discrimination computed over the period 1992-1995 and merged with firms in 1996-2001. Model 3 adds interactions and controls for municipalities to the regression. Model 4 contains the same specification of Model 3, but excludes most gender segregated firms in terms of occupational composition.

a specification that allows the labour supply to the firm to differ across a finer set of dimensions.[39] For the same reason, we interact all industry fixed effects with local labour market fixed effects. Finally, we also add fixed effects for the smallest geographical entities available in the data, *i.e.* municipalities. In Model 4 we further limit the potential problem of measurement error in gender differences in firm wage policies, by excluding from the estimation sample firms in which the *gender dissimilarity index* in job titles is above the 75th percentile. This restriction implies the exclusion of all firms in which this index exceeds a level of around 0.45 and an obvious reduction in sample size.[40]

As can be noticed from the table, when estimating these two last specifications, the overall fit of the regression improves. Moreover, the main coefficients of interest remain quite stable and significant, despite the substantial growth in the number of dependent variables. For example, in Model 4, where highly segregated firms are excluded from the sample, the point estimate of the effect of the presence of women in the management even grows, as well as the confidence level associated to this coefficient.

Overall, the last two columns of Table 4 provide the most solid evidence in support of the hypothesis that taste-based discrimination is indeed one of the drivers of gender differences in firms' pay premiums. However, these results also suggest that other mechanisms, such as gender differences in compensating wage differentials, may be relevant in determining this gap and they should be controlled for and further tested whenever possible. Finally, these results show that both, the female share of employment within the firm and the presence of women at the top of the occupational hierarchy are coherent proxies for discriminatory preferences within a workplace, as they are robust to a formal test on their validity.

6 Conclusions

In this paper, we have shown that a simple static model of taste-based discrimination in monopsonistic labour markets can provide a coherent framework to interpret the gender

[39] Refer to Section 2.3 for a more formal discussion on the main assumptions behind this model specification and on its ability to control for monopsonistic discrimination.

[40] See Section 2.4 for a definition of the dissimilarity index and for a more formal discussion on the rationale behind this inclusion restriction.

wage gap in firms' pay policies. This component of the earning differential can be recovered from an AKM regression model (Abowd et al. [1999]), and it has been recently documented by Card et al. [2016] using Portuguese data. We have provided conditions under which this residual component of the gender wage gap can be attributed to taste-based discrimination. Moreover, we have proposed empirical methods that, under reasonable assumptions, allow to recover this parameter controlling for potential confounding factors, most notably gender differences in the labour supply to the firm and within workplaces occupational segregation.

We have applied the proposed methodology to Italian data, showing that women employed in the right tail of most discriminatory firms suffer additional wage losses with respect to men of between 5% and even more than 10%, depending on the definition of the treatment effect. Moreover, we have shown that this component of the gender wage gap can not be explained by gender-specific returns to observable characteristics in a traditional cross-sectional setting.

By documenting a positive relationship between the residual component of the gender wage gap, as identified through our regression framework, and traditional proxies associated to discriminatory preferences, we have shown that our main findings are indeed coherent with a model of taste-based discrimination. In particular, we have tested the hypothesis that gender differences in firms' premiums are negatively correlated with the presence of women at the top of the firms' hierarchy and with the share of female employment within workplaces. We have found strong evidence supporting both results, but we have shown that also other mechanisms, such as compensating wage differentials, are probably driving the overall firm-specific residual gender gap.

The above results are not only coherent with our theoretical framework, but they also support the validity of two firms characteristics usually advocated as potential measures of employers' preferences against women. For this reason, we believe that the methods described in this paper can provide interesting insights to future research on the implication of Becker's theory, as they allow to conduct formal tests on the validity of proxy variables measuring discrimination. More generally, we believe that future research on the effectiveness of affirmative action policies could also potentially benefit from the application of the model presented here, improving our understanding of the most important mechanisms driving discriminatory differences in wages also in other contexts.

Appendix

A Distribution of the Firm-Specific Gender Wage Gap

In this Appendix, we provide descriptive statistics on the residual firm-specific gender wage gap defined in equation (6), focusing in particular on how accounting for the model's main independent variables affects its distribution across local labour markets and gender groups.

Figure A.1 shows how the inclusion of the main controls for size, industry and occupation discussed in Section 5.2 affects the gender pay gap in firms' residuals. The left panel of the figure shows a map of the local labour markets included in the analysis, in which districts displayed in a darker colour have a relatively higher median gender wage gap in firms' residuals. In the right panel, we report the same statistic adjusted for differences in the size, sector and occupational composition of firms within each local labour market. To ease the interpretation of the quantities, we have expressed the raw and conditional gender wage gaps as differences from the lowest observed raw gender wage gap in the sample.

In comparing the left and right panels of Figure A.1, notice that the relative ranking of districts in which women do relatively better changes with the inclusion of independent variables. Districts located toward the South-East tend to show lower gender differences than districts located in the more Central and Northern areas, even if there are some exceptions to this trend.[41] Finally, notice that the range from the highest to the lowest median gender wage gap reduces when considering the conditional distribution instead of the raw difference in residuals, as does the overall size of the gender gap.

Figure A.2 shows kernel density estimates of the residual derived from equation (6), as es-

[41] Veneto is relatively flat toward the South. The South-Eastern borders are on the sea, while South-Western borders are shared with the flat areas of two important manufacturing regions of Italy, *i.e.* Lombardia and Emilia-Romagna. Veneto is instead more mountainous in its more Central and Northern parts, even if important manufacturing district are present also in some of these areas.

Figure A.1: **Firm-Specific Gender Wage Gap by Local Labour Market**

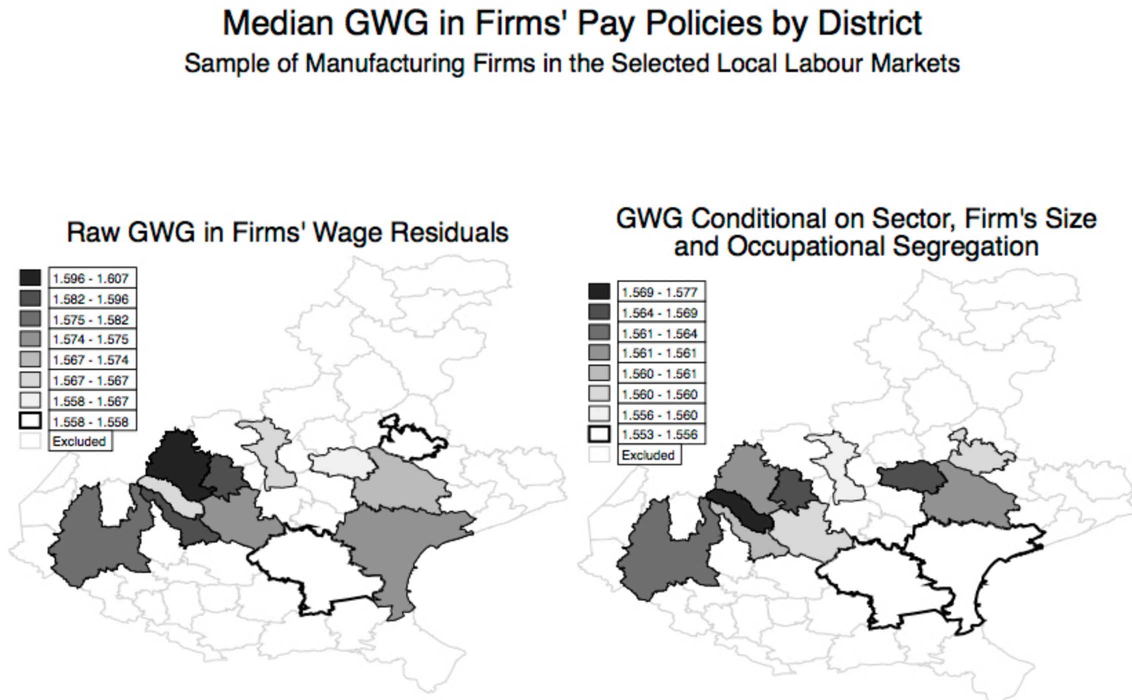

The median is computed over firms, without weighting for their size. The conditional distribution of the gender wage gap is obtained by regressing the raw gender difference in firm wage policies on industry dummies (two-digits), dummies for each quartile of the firm size distribution (average yearly total days worked within the firm) and on gender differences in the proportions of occupational composition (five occupations for each two-digit sector). The conditional gender wage gap is then defined as the residual of the regression. The residual and the raw gender wage gap are expressed as a difference from the lowest observed raw gap in the sample.

Figure A.2: **Distribution of the Residual Gender Wage Gap in Firms' Premiums**

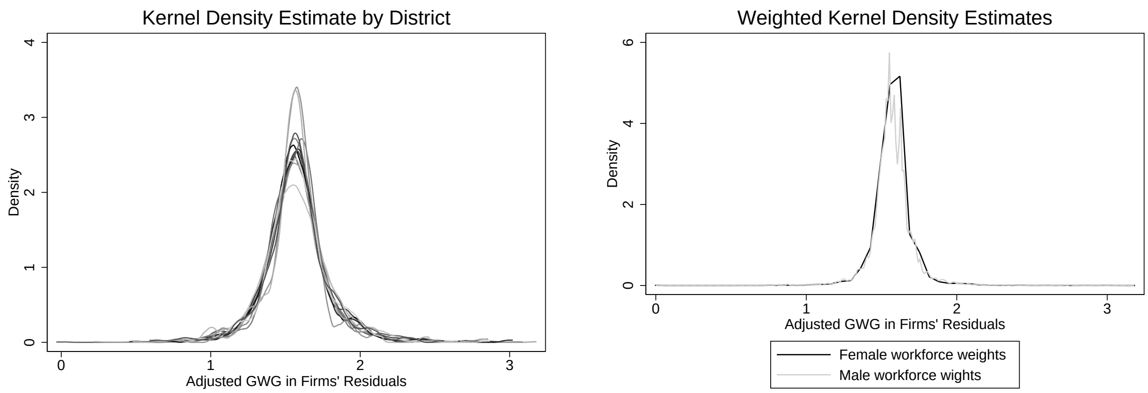

The residual gender wage gap is estimated by regressing the raw gender difference in firm wage policies on district fixed effects, industry dummies (two-digits), dummies for each quartile of the firm size distribution (average yearly total days worked within the firm) and on gender differences in the proportions of occupational composition (five occupations for each two-digit sector). The residual is expressed as a difference from its lowest observed value in the sample.

timated including the following independent variables: local labour market fixed effects, two-digit sectors fixed effects, dummies for each quartile of firms' size and within-firm gender differences in the proportion of workers in each job title. In this case, we have expressed this residual as a difference from its lowest observed value in the sample.

In the right panel of Figure A.2, the kernel densities are estimated separately by local labour market and considering only one observation per firm. Given that district fixed effects are included, the residual is approximatively normal and its distribution is quite similar across local labour markets. Considering this unweighted distribution of the variable, the standard deviation is 0.225 and the growth of the gender wage gap from the 5th to the 95th percentile is of about 70%.

These numbers change once that firms' size is accounted for. The right panel of Figure A.2 shows the distribution of the regression residual weighted for the (gender-specific) size of establishments. The dispersion in this residual is smaller when computed over the workforce.[42] Another interesting feature is that women are relatively more likely to work at more discriminatory firms, given that their distribution is shifted to right with

[42] In this case, the variance of the firm-specific residual gender wage gap becomes 0.134, and the growth of the gap from the 5th to the 95th percentile of the distribution is of only around 37%.

respect to that of men. A possible hypothesis on such pattern is that it could be driven by compensating wage differentials, which, due to hedonic considerations, would make discriminatory workplaces more attractive to women despite the wage penalty associated with them. However, such hypothesis should not be based solely on the evidence of Figure A.2, given that the result is quite preliminary, subject to alternative mechanisms and prone to measurement error issues. Section 5.4 discusses a more rigorous test on the relevance of compensating wage differentials considerations, using a parametrization based on the likelihood with which firms provide flexible labour market contracts.

References

Abowd, J., F. Kramarz, and D. Margolis (1999): "High Wage Workers and High Wage Firms," *Econometrica*, 67, 251–333.

Abowd, J. M., R. H. Creecy, and F. Kramarz (2002): "Computing Person and Firm Effects Using Linked Longitudinal Employer-Employee Data," Longitudinal Employer-Household Dynamics Technical Papers 2002-06, Center for Economic Studies, U.S. Census Bureau.

Andrews, M. J., L. Gill, T. Schank, and R. Upward (2008): "High wage workers and low wage firms: negative assortative matching or limited mobility bias?" *Journal of the Royal Statistical Society Series A*, 171, 673–697.

Arulamplam, W., A. Booth, and M. Bryan (2006): "Is There a Glass Ceiling in Europe? Exploring the Gender Pay Gap Across the Wage Distribution," *Industrial and Labor Relations Review*, 62(2), 163–186.

Ashenfelter, O. and T. Hannan (1986): "Sex discrimination and product market competition: The case of the banking industry," *The Quarterly Journal of Economics*, 149–173.

Barth, E. and H. Dale-Olsen (2009): "Monopsonistic discrimination, worker turnover, and the gender wage gap," *Labour Economics*, 16, 589–597.

Becker, G. (1957): *The economics of discrimination*, University of Chicago Press, Chicago.

Black, D. A. (1995): "Discrimination in an equilibrium search model," *Journal of Labor Economics*, 309–334.

Black, S. E. and E. Brainerd (2004): "Importing equality? The impact of globalization on gender discrimination," *Industrial & Labor Relations Review*, 57, 540–559.

Blau, F. D. and L. M. Kahn (2017): "The Gender-Wage Gap: Extent, Trends, and Explanations," *Journal of Economic Literature*, forthcoming.

Boal, W. M. and M. R. Ransom (1997): “Monopsony in the labor market,” *Journal of Economic Literature*, 35, 86–112.

Card, D., A. R. Cardoso, J. Heining, and P. Kline (2017): “Firms and Labor Market Inequality: Evidence and Some Theory,” *Journal of Labour Economics*, forthcoming.

Card, D., A. R. Cardoso, and P. Kline (2016): “Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women,” *The Quarterly Journal of Economics*, 131, 633–686.

Charles, K. K. and J. Guryan (2008): “Prejudice and Wages: An Empirical Assessment of Becker’s The Economics of Discrimination,” *Journal of Political Economy*, 116, 773–809.

Depew, B. and T. A. Srensen (2013): “The elasticity of labor supply to the firm over the business cycle,” *Labour Economics*, 24, 196–204.

Flabbi, L., M. Macis, A. Moro, and F. Schivardi (2014): “Do Female Executives Make a Difference? The Impact of Female Leadership on Gender Gaps and Firm Performance,” IZA Discussion Papers 8602, Institute for the Study of Labor (IZA).

Graham, B. S. (2017): “Identifying and Estimating Neighborhood Effects,” *Journal of Economic Literature*, forthcoming.

Hellerstein, J. K., D. Neumark, and K. R. Troske (2002): “Market Forces and Sex Discrimination,” *Journal of Human Resources*, 37, 353–380.

Heyman, F., H. Svaleryd, and J. Vlachos (2013): “Competition, takeovers, and gender discrimination,” *Industrial & Labor Relations Review*, 66, 409–432.

Hirsch, B., T. Schank, and C. Schnabel (2010): “Differences in labor supply to monopsonistic firms and the gender pay gap: An empirical analysis using linked employer-employee data from Germany,” *Journal of Labor Economics*, 28, 291–330.

Kawaguchi, D. (2007): “A market test for sex discrimination: Evidence from Japanese firm-level panel data,” *International Journal of Industrial Organization*, 25, 441–460.

Lorenzini, F. (2005): “Distretti industriali e sistemi locali del lavoro 2001,” in *8 Censimento generale dell’industria e dei servizi*, ISTAT.

Manning, A. (2003): *Monopsony in motion: Imperfect competition in labor markets*, Princeton University Press.

Matsudaira, J. D. (2014): “Monopsony in the Low-Wage Labor Market? Evidence from Minimum Nurse Staffing Regulations,” *The Review of Economics and Statistics*, 96, 92–102.

Muehlemann, S., P. Ryan, and S. C. Wolter (2013): “Monopsony Power, Pay Structure, and Training,” *ILR Review*, 66, 1097–1114.

Oaxaca, R. (1973): “Male-Female Wage Differentials in Urban Labor Markets,” *International Economic Review*, 14(3), 693–709.

Oi, W. Y. and T. L. Idson (1999): "Firm size and wages," in *Handbook of Labor Economics*, ed. by O. Ashenfelter and D. Card, Elsevier, vol. 3 of *Handbook of Labor Economics*, chap. 33, 2165–2214.

Ransom, M. R. and R. L. Oaxaca (2010): "New market power models and sex differences in pay," *Journal of Labor Economics*, 28, 267–289.

Ransom, M. R. and D. P. Sims (2010): "Estimating the Firm's Labor Supply Curve in a "New Monopsony" Framework: Schoolteachers in Missouri," *Journal of Labor Economics*, 28, 331–355.

Robinson, J. (1933): *The economics of imperfect competition*, Macmillan, London.

Sulis, G. (2011): "What can monopsony explain of the gender wage differential in Italy?" *International Journal of Manpower*, 32, 446–470.

Webber, D. A. (2015): "Firm market power and the earnings distribution," *Labour Economics*, 35, 123–134.

Weber, A. and C. Zulehner (2014): "Competition And Gender Prejudice: Are Discriminatory Employers Doomed To Fail?" *Journal of the European Economic Association*, 12, 492–521.

Chapter II

Collective Bargaining and the Evolution of Wage Inequality in Italy*

Abstract

This paper studies the evolution of Italian male wage inequality over a two-decade period, showing that pay differences have increased since the mid-1980s at a relatively fast pace. By accounting for worker and firm fixed effects, it is shown that observed and unobserved heterogeneity of the workforce have been major determinants of increased wage dispersion, while variability in firm wage policies has declined over time. The growth in wage dispersion has entirely occurred between *livelli di inquadramento*, *i.e.* job titles defined by national and industry-wide collective bargaining institutions, and for which specific minimum wages apply. These results suggest that the underlying market forces determining wage inequality have been largely channelled into the tight tracks set by the country's fairly centralized system of industrial relations.

JEL Codes: J00, J5, J31, J40.

Keywords: Wage Inequality; Collective Bargaining; Firm Wage Policy; Two-Way Fixed Effects; Matched Employer-Employee Data.

*This article is the product of a joint work with Francesco Devicienti (University of Torino) and Agata Maida (University of Milano). I would like to thank David Card, Lia Pacelli, participants at the 2016 European Association of Labour Economists Conference, 2016 Conference of the European Society of Population Economics, and Collegio Carlo Alberto workshops for useful comments and feedbacks.

1 Introduction

Wage inequalities have risen in most Western Countries during the last decades of the past century. Several hypotheses have been put forward to rationalize this secular trend.[1] Some authors have pointed out that technological progress is largely responsible for the increased wage dispersion. A commonly held view is that advances in the production process may have led to a growth in the demand for skills that, due to demographic and schooling developments, outpaced their supply, eventually resulting in an increase of the returns to unevenly distributed workers' characteristics (e.g. Katz and Murphy [1992]).[2] Other theories state that changes in labour market institutions, such as declining minimum wages and union strength (e.g. Di Nardo et al. [1996]), or changes in social norms (e.g Piketty and Saez [2003]) are the main drivers behind the observed secular rise in wage differentials.

More recently, several studies analysing matched worker-firm databases have pointed out that one important component of pay inequality is represented by differences in the wage policies between observationally similar firms (e.g. Abowd et al. [1999]). Card et al. [2013] have shown that firm-specific components of the wage variance can explain up to one fourth of the inequality growth occurred in West Germany between the late-1980s and the beginning of the new century. Moreover, they show that the sorting of high skilled workers into high-paying firms can account for another 30% of the variance increase. They link the rise in the dispersion of firms' wage premiums to the changes that have occurred in the wage bargaining system since the early 1990s. Indeed, during these years German firms often exploited the possibility of *opting-out* from national contractual agreements, deviating from collective bargaining provisions and resorting to establishment-level negotiations. Dustmann et al. [2014] argue that this decentralization in the wage setting process has allowed to cut unit labour costs and to improve international competitiveness, fostering the German economic growth observed in the last decade.

[1] For a recent review of the vast literature on wage inequality see Acemoglu and Autor [2011].

[2] A more nuanced hypothesis, also related to technological progress, is that recent advances in the production process may have modified the demand for routine versus non-routine based occupations, increasing the polarization of the wage structure (see for example Autor et al. [2008]).

In this paper, we apply the methodology of Card et al. [2013] on Italian matched employer-employee data from administrative sources, which cover the entire population of private-sector workers and firms in the Veneto region. Considering the male sample only, we study the evolution over the period 1982-2001 of the following components of pay dispersion: time-varying characteristics of the workforce, time-constant individual characteristics, firm-specific wage premiums, along with the contribution arising from the the correlation between each of these components. Moreover, we apply a related variance decomposition method, in order to test whether the growth in wage inequality has occurred mostly within or between the fine job title categories defined by the country's collective bargaining institutions.

An in-depth analysis of the Italian case is interesting in itself, but it also offers insights for evaluating the relevance of the various theories rationalizing the secular growth in wage inequality experienced elsewhere. The Italian labour market is characterized by sector-wide collective bargaining and, like other Western Countries, the Italian economy has been exposed to international competition and has experienced the challenges posed by the introduction of new technologies. Moreover, the manufacturing sector is very large in Italy, and particularly in Veneto, a feature that makes its economy quite similar to the German one.

Our analysis shows that, during the overall period considered, Italian pay dispersion has grown at a similar pace than the one documented by Card et al. [2013] for Germany. However, Italy did not experience any growth in firms' wage premiums dispersion, given that the variance of this wage component is even declining over time. Considering that Italian wage setting mechanisms are highly centralized at the sector-wide level and have not undergone the same renewal processes characterizing the German labour market during the 1990s, our results suggest that the amount of wage flexibility granted to employers by such system has not grown over time. That is to say, Italian firms have been unable to opt-out, or diverge in any other significant way, from the wage dynamics settled within the relevant industry-wide collective agreements.

We find that a large proportion of the growth in earning dispersion over the entire period considered is due to raising heterogeneity in the *portable* component of a worker's pay, namely the part of the wage attributable to individual-specific characteristics equally rewarded across employers. In principle, a growing contribution of workers' heterogeneity to the total wage variance may simply reflect the underlying dynamics of supply and demand factors. However, we show that in practice this component of inequality is closely linked to the wage pay scales and seniority wage increments bargained at the industry level by the main union confederations and employers' associations. Hence, we interpret the finding of rising workers' heterogeneity as yet another outcome induced by the Italian system of industrial relations, which seems to impose significant constraints on wage dynamics.

To substantiate our claim, we divide the variance of (log) wages and of workers' *portable* pay components into a *within* and a *between* job titles part. Job titles (called *livelli di inquadramento* in Italian) are occupations defined by the relevant sectoral collective agreements, for which a specific minimum wage applies regardless of a worker's union membership. We find that the growth in the *between*-variance component virtually explains the entire inequality trend observed in the data. To the best of our knowledge, this evidence has never been so extensively documented before, partly owing to data limitations in past research on Italian wage inequality.[3]

Our analysis shows also that another important component of the growth in wage inequality has been increased positive sorting between firms' pay premiums and the human capital of the workforce. Despite a low level of correlation between these two components in each time period, a clear increasing tendency emerges from our estimates. Although we were unable to present conclusive evidence on the determinants of such trend, it is tempting to associate at least part of the growth in assortative matching to the general labour market deregulations experienced by Italy since the mid-1980s.

[3] A nice feature of our data is that they contain job-title information for every worker in the sample. The datasets used in previous studies on Italy, e.g. the Bank of Italy's Survey on Household Income and Wealth used by Manacorda [2004], do not contain individual-level information on *livelli di inquadramento* or, when they do, cover only a specific sector (metalworkers) and area (the province of Milan), e.g. Erickson and Ichino [1995]. Using Portuguese data, Torres et al. [2013] find that job title fixed effects in wage regressions explain around 10% of the total wage variance. However, they do not explicitly focus on the contribution of these effects in explaining the *evolution* of inequality in Portugal.

Our paper is connected to a (moderately-sized) literature that has previously studied the Italian wage distribution and its changes over time, using either social security data (see, among others, Cappellari [2004]) or earnings data collected in household surveys (e.g. Manacorda [2004]). Nevertheless, our use of the decomposition methodology proposed by Card et al. [2013] is entirely new in the Italian context.[4] To the best of our knowledge, our finding of a negative contribution of the firm-specific wage premiums to the trends in overall wage inequality is also novel, and potentially interesting, beyond Italy.

The paper is organized as follows. In the next section, we review the existing literature on the evolution of Italian wage inequality, providing a brief institutional framework. Section 3 describes the database and preliminary evidences on pay dispersion. Section 4 reviews the main econometric model employed in the analysis and discusses its assumptions. Section 5 presents the main results. Finally, Section 6 contains the concluding remarks.

2 Institutional Context and Related Literature on Italian Wage Dispersion

During the years considered in this study (1982-2001), and largely still today, Italy has been characterized by a wage setting mechanism fairly centralized at the sector-wide national level. Collective contracts are *de-facto* binding for all employers and all workers, irrespective of union membership. Such agreements are signed (typically every two years) by the major trade unions and employers' associations at the industry-wide level. Each contract regulates specific job titles (*livelli di inquadramento*) and the contractual minimum wages that is to apply for each of them. There are no *opting-out clauses*. That is to say, firms cannot decide to resort to firm-level contractual agreements derogating to the wage standards settled at the sectoral level.[5] Regional- or firm-level agreements can only distribute *top-up* wage components, typically related to indicators of profitability or

[4]The two-way fixed effects model has been estimated on Italian data in different contexts (*e.g.* Iranzo et al. [2008] and Flabbi et al. [2014]). However, none of these studies has focused on wage inequality and on its evolution over time.

[5]Only in 2011, well beyond the time period studied in this paper, an agreement signed by trade unions and employers' organizations has attempted to widen the scope for derogation of firm-level contracts with respect to sectoral bargaining. Nevertheless, reforms in the wage setting institutions remain an actively debated policy topic in the Italian context.

productivity.

In 1993 a major reform of collective bargaining was approved, in order to achieve the following main objectives:[6] (*i*) coordination across industries and moderation on wage growth to achieve low inflation targets; (*ii*) growth of regional differences in wages to adapt them better to the heterogeneous cost of living and labour market conditions at the local level; (*iii*) distribution of premiums related to performance (*on top* of the sectoral minimums) and negotiation of some other contractual provisions not related to compensation at the firm-level. This reform resulted in an increase of geographical differences in top-up components of negotiated wages. However, Devicienti et al. [2008] find that overall the amount of flexibility in bargaining agreements introduced by the 1993 reform has been quite limited.[7]

The years under study were characterized by a relatively persistent economic growth, even if periods of turmoil were not completely absent. The left panel of Figure B.1 (in the Appendix) shows that the GDP in Italy and in Veneto grew from the early 1980s until 1991, except for a short period of stagnation, which lasted until 1984. The years from 1991 to 1993 were instead characterized by an economic recession, which was followed by a recovery phase during which North-Eastern Italy started outperforming the national economic growth level. The right panel of Figure B.1 compares the evolution of male and female absolute levels of employment in the Veneto private sector.[8] It shows that, during the overall period under study, men's employment level grew by slightly less than 20%, while the same figure is almost 50% for women. Such large gender differences in the evolution of participation are one of the main reasons that led us to study the men's and women's sample separately, focusing the present analysis on male wage inequality only.[9]

Figure B.2 provides an overview of the long-run evolution of real gross *weekly* wages, computed from the social security records of male private sector workers in Veneto. The

[6]Casadio [2003] provides a detailed review of the content of this reform. Moreover, a detailed institutional background on the pre-1993 context is given by Erickson and Ichino [1995].

[7]Using a sample covering around 60% of national private-sector contracts, these authors show that the average share of all top-up components over total wages increased from around 18% during the mid-1980s, to only 22% by the end of the 1990s.

[8]This figure is constructed from the social security records analysed throughout this paper.

[9]Results obtained in the female sample are presented in a working-paper version of this article, see Devicienti et al. [2016].

average pay, which is reported in the left panel of the figure, increased up until the early 1990s, when a phase of economic crisis begun and wage growth became persistently flat. The right panel of Figure B.2 shows that inequalities, as measured by the standard deviation of log weekly wages, declined sharply until around 1983. Previous research has attributed this remarkable trend to the strong compressing effects of the *Scala Mobile* (see, in particular, the analysis in Manacorda [2004], based on household survey data).[10] The *Scala Mobile* was a cost-of-living allowance added quarterly to the bargained contractual minimum wages, an institution provided from the 1970s until 1993, but which was majorly reformed in 1984 and then through a referendum in 1985. Since this wage-adjustment mechanism had been particularly disadvantageous for more qualified white-collars and skilled workers, from 1987 on, in order to further mitigate its egalitarian effects, most nation-wide collective bargaining agreements improved the compensations associated to the qualifications embedded in each *livello di inquadramento*, widening the gaps in the minimum wages stipulated for each of these job titles.[11] The right panel of Figure B.2 shows indeed that the period 1982-2001, the one on which we focus our analysis, is instead characterized by a very persistent growth in pay inequality.

Other papers have studied the Italian wage inequality from the early 1980s onward, documenting a growth in pay dispersion. The main data sources that have been used are the *Bank of Italy Survey on Income and Wealth* (see Brandolini et al. [2002], Manacorda [2004] and Naticchioni and Ricci [2009] among others) and the *Worker History Italian Panel* (WHIP), or similar administrative data containing samples of the private sector of the entire Italian territory (see, for example, Devicienti and Borgarello [2001], Cappellari [2004] and Cappellari and Leonardi [2016]). However, only a few papers have specifically looked at the role played by collective bargaining institutions in the evolution of wage inequality. In particular, Dell'Aringa and Lucifora [1994] and Erickson and Ichino [1995], using a sample of metal-mechanical workers in the metropolitan area of Milan, argue that

[10] Leonardi et al. [2015] further investigate this issue, documenting the presence of substantial wage penalties for high skilled workers employed in firms more affected by the Scala Mobile during the 1976-1982 period.

[11] This general tendency is often highlighted by industrial relations reports of the time (*e.g.*, CESOS [1989]).

the centralized system of industrial relations plays a pivotal role in determining the Italian wage structure. Beyond their narrower coverage, the data employed in these studies have other limitations that impede the estimation of a two-way fixed effect model, and of the related wage variance decomposition (e.g., an insufficient coverage of each firms' entire workforce or a limited panel dimension).

Differently from the papers reviewed above, in the present analysis we directly look at the contributions of firm wage policies, worker heterogeneity and assortative matching to the evolution of wage inequality. Moreover, using information on the *livelli di inquadramento* (*i.e.* job titles) for the entire workforce, we are able to uncover the role played by collective bargaining institutions from a more comprehensive perspective.

3 Data and Preliminary Evidences on Inequality

3.1 Database and Descriptive Statistics

The Veneto Working Histories (VWH) database, which is studied here, contains earnings data from social security records for dependent workers of the private sector in the Veneto region. The database contains the *population* of private sector firms whose headquarters are located in Veneto, and the *population* of their employees.[12] In order to analyse a sample of workers more homogeneous and consistent across time, we have divided the data by gender and, throughout this paper, we discuss only results obtained among men.[13]

The information contained in social security records, even if not rich on some workers' characteristics, is highly accurate and reliable, since employers are obliged to report such information correctly by law. A limitation of these data, dictated by their current availability, is that they allow to estimate a two-way fixed effect regression model for only one region of the country.[14] Nevertheless, the main inequality trends observed in the VWH

[12]Workers of these firms are followed if they continue working in a private-sector establishment outside the Veneto region, but they are not followed if they move to the public-sector. Moreover, for firms outside Veneto information on the entire workforce is not available.

[13]Results for women, presented in a working paper version of this article (Devicienti et al. [2016]), show general trends that are similar to the ones observed among men. However, there are also some differences in the results, most likely driven by gender-specific employment dynamics (see Figure B.1) and by the incidence of part-time work, which is much higher and growing faster in the female sample.

[14]Currently available country-level longitudinal matched employer-employee data are not suitable for our aims. For example, the Worker History Italian Panel is a 1:90 sample of dependent workers, entailing

data are very similar to those observable in other national level samples of social security records (such as WHIP). There are other circumstances that make Veneto a particularly informative case-study. First, this region has a well-developed manufacturing sector and is fairly large, given that its population amounted to 4.5 millions in 2001 and its economy represented around 11% of the national GDP in the same year. Finally, a second and more subtle advantage of analysing a single region of Italy is given by the fact that the approval of the 1993 industrial relations reform, which introduced more flexibility at the regional level in order to better link wages to local market conditions, could induce an over-estimation of trends in wage inequality when studying national-level data, given the difficulty in controlling for genuine adjustments of wages to the regional market conditions.[15]

The VWH database contains information on wages from 1975 until 2001,[16] but we consider only the last two decades of the data. In particular, in the rest of the paper we study the years from 1982 to 2001, since our main purpose is to shed light on the determinants of the inequality growth, which takes place during this most recent period. Moreover, information on days worked is reported only starting from 1982 and we choose log gross *daily* wages, adjusted to the 2003 level, as the unit of measurement for earnings, since this is the most precise measure controlling for time worked.[17]

We have taken a number of steps that are relatively standard in the literature using similar data. First, for each employee with multiple jobs during the same year, we have selected the most representative spell in terms of months, weeks and days worked, resorting to total earnings to break the few remaining ties. Second, we have excluded from the sample all spells shorter than approximately four months (16 weeks) and, finally, we have trimmed wages at the 1st and 99th percentiles calculated over a six-year period.

that for many firms only a handful of workers are observed. The only other employer-employee data source that has been analysed in the literature (*i.e.* the one studied by Iranzo et al. [2008]) allows to observe the entire workforce of each firm, but only relatively large firms are included in the sample.

[15] On this topic, Devicienti et al. [2008] document a tenuous *resurrection of the Italian wage curve* after 1993, mostly driven by greater regional differences in wages.

[16] Like for other social security data, the information on pay is gross of taxes and inclusive of all cash benefits, but it excludes all in-kind benefits.

[17] Other available alternatives (weekly or monthly wages) are less precise since, by the law, employers have to report all weeks and months during which an employee has worked *at least one day*.

Table 1: **Summary Statistics (Mean and St. Dev.) by Period**

Period	**1982-1987**	**1984-1989**	**1988-1993**	**1992-1997**	**1996-2001**
Log daily wages	4.782	4.801	4.855	4.866	4.874
St. Dev.	0.287	0.303	0.337	0.350	0.362
Age	36.76	36.37	35.94	35.83	35.81
St. Dev.	11.07	11.04	10.91	10.40	9.85
Firms' workers	7.884	7.392	7.031	7.500	7.419
St. Dev.	55.30	48.77	39.05	52.09	49.023
Tenure	5.072	5.552	6.013	6.451	6.479
St. Dev.	3.631	4.325	5.494	6.252	6.823
Proportions					
Part Time	0.002	0.004	0.007	0.011	0.016
Apprentice	0.016	0.020	0.025	0.025	0.035
Blue Collar	0.730	0.729	0.723	0.723	0.708
White Collar	0.247	0.243	0.242	0.244	0.250
Manager	0.007	0.007	0.008	0.007	0.007
Primary Sect.	0.043	0.045	0.045	0.044	0.042
Secondary Sect.	0.626	0.631	0.648	0.651	0.663
Tertiary Sect.	0.331	0.324	0.307	0.305	0.295
Total Workers	**698,378**	**724,448**	**753,753**	**777,019**	**845,984**
Total Firms	**64,972**	**72,605**	**80,159**	**80,572**	**85,104**

The sample is composed of firms located in Veneto belonging to the largest connected set. Part-time contracts have been introduced only since 1985. Tenure is censored at 1975.

To estimate the two-way fixed effect model of Abowd et al. [1999] we have divided the 1982-2001 years of data into five, partially overlapping, six-years panels. Moreover, we have included only observations belonging to the *largest connected set*.[18] and all results are computed including only firms of Veneto.[19]

Table 1 contains descriptive statistics for each of the five panels that we have constructed. The composition of the sample is quite homogeneous across periods, even if some noticeable trends emerge. Given that civil servants are excluded from the social security archives and that Veneto is a manufacturing-oriented economy, the secondary sector is relatively large and this pattern is reflected in the occupational composition of the sample (the majority of individuals are blue-collars). The (non-weighted) average firms' size,

[18]This is a set of workers and firms connected by employees' mobility. This restriction implies the loss of an extremely small proportion of observations (around 1-2 %), and is usually applied to simplify estimations. See Abowd et al. [2002] for a discussion.

[19]However, we have included also employment spells outside this region in the estimation sample of the two-way fixed effect model. The rationale of these choices is discussed in the section providing the details of our econometric method.

Figure 1: **Dispersion of Log Daily Wages by Year**

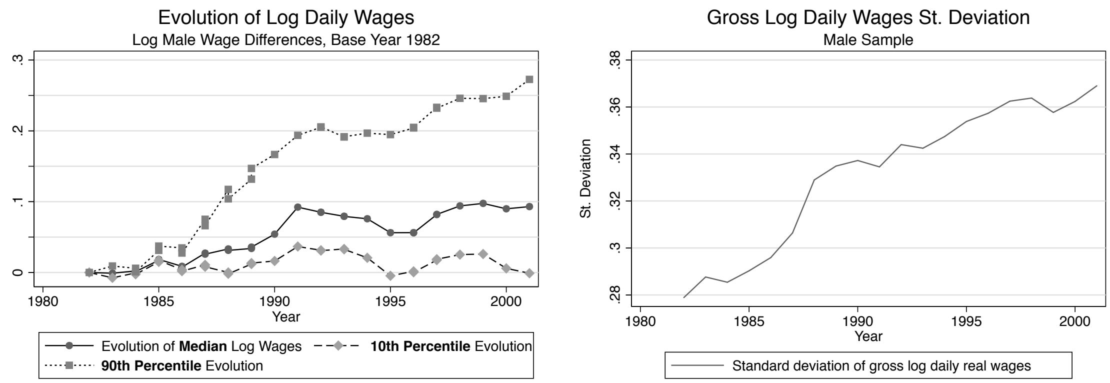

measured by the number of employees continuously working for at least six months in a year, tends to be quite small and it is slowly decreasing over time. Average age is slightly decreasing over time, while the opposite is true for tenure. However, in this latter case the result is driven by the fact that tenure is left-censored at the year 1975.[20] Finally, the percentage of part time contracts is relatively low and it grows over time, a tendency attributable to the fact that such contracts have been introduced in the Italian legislation only since 1985.

Table 3.1 shows also that real wages have been quite flat during the overall period considered, while their dispersion, as measured by the standard deviation, steadily increases. In the next paragraph, we present a more accurate description of this trend.

3.2 Preliminary Evidences on Inequality

Figure 1 describes, in the left panel, the evolution of log daily wages at the 10th, 50th and 90th percentiles of the earning distribution and, in the right panel, the evolution of their standard deviation. It can be noticed that the standard deviation, the 50th-10th and 90th-50th wage percentile ratios have all increased. Another evidence emerging from the table is the very slow, and even slightly negative growth of wage levels at the bottom of the distribution. In particular, men's earnings at the 10th percentile have remained stable over the whole period, median wages have risen by almost 10% and the 90th per-

[20]To correct for this problem, in the empirical analysis we control for tenure by adding dummy variables for the first six years of tenure, leaving higher seniority levels as the reference category.

centile of the pay distribution has steadily risen by almost 30% in real terms, except for a stagnating period during the early- and mid-1990s.

In the left panels of Figure 2 we test the predictive performance of a series of log-linear conditional wage models. To construct this figure, we have run year-by-year OLS regressions on the workers of firms located in Veneto, using different sets of controls. The highest line of the left-panel graph represents the unconditional log wage variance. The other lines represent the root mean squared error (RMSE) of year-specific regression models. In each model, we have used the same set of base-line covariates, namely: a quadratic in age, occupation dummies, tenure dummies, log of firm size (number of employees), around thirty sector fixed effects, national industry-wide collective contract fixed effects, a set of interactions (age with occupation and age with tenure). In addition to these covariates, each regression model is fully saturated for one of the following categories: (1) job titles (*livelli di inquadramento*), (2) firms or (3) both.[21]

The RMSE provides a measure of the performance of each model in explaining total wage variation. In general, the trend in residual wage variance is fairly flat, while the total pay variance shows a clear increasing pattern. This is a preliminary evidence that workforce composition and returns to its characteristics do a good job in explaining the rise in wage dispersion over time, and are becoming increasingly relevant over time. Obviously, none of the estimated models reported here does control for constant unobserved individual characteristics. Nevertheless we can see that a fairly small proportion of the unconditional wage variation remains unexplained, especially when we estimate a model fully saturated for job titles and firms. Firm fixed effect explain a greater proportion of wage variation than job title fixed effects. However, when focusing on the evolution of the RMSE across time, the same pattern does not hold.

In order to better compare the evolution of the relative performance of each of the three regression specifications, in the right panel of Figure 2 we normalize each year-specific

[21] National industry-wide collective contract fixed effects are not collinear with job title fixed effects, since the latter are specific occupations (usually between five and ten) defined by the former. Instead, firm fixed effects are collinear with sector fixed effects and, typically at least, also with industry-wide contracts fixed effects. The procedure adopted in constructing job title and collective contract fixed effects is discussed in more detail in Section 5.3.

Figure 2: **Wage St. Dev. and RMSE from Alternative Wage Models**

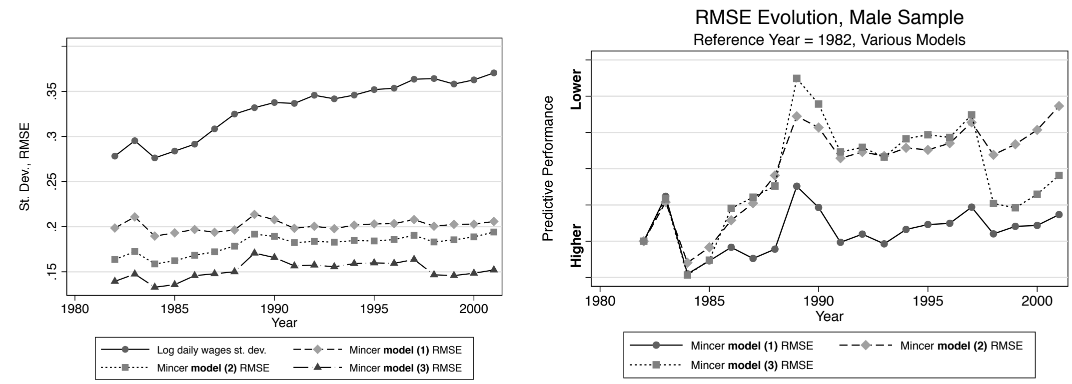

Baseline controls: age (quadratic), tenure dummies, four qualification dummies, log of employees number, sector fixed effects, national industry-wide collective contract fixed effects.

Models' definition: *(1)* job title (*livello di inquadramento*) fixed effects; *(2)* firm fixed effects; *(3)* fully saturated fixed effects for job titles and firms.

Note: national collective contracts vary within and across sectors, and might be not homogeneous across years. *Livelli di inquadramento* are occupational positions determined by each national collective contract, and are not homogeneous across years.

RMSE to the 1982 level of the corresponding model.[22] A clear pattern emerges, as over time the explanatory power of fixed effects for job titles gains importance with respect to the models where firm effects are controlled for. Since *livelli di inquadramento* are defined by sector-wide collective agreement, and a particular minimum wage is set for each of these occupational positions, we interpret the right panel of Figure 2 as a preliminary evidence of the importance of collective bargaining in shaping the evolution of the pay distribution. In Section 5.3, employing a more informative regression framework, we analyse this point in more detail.

4 Econometric Methodology

The contributions of firm-specific, time-constant and time-varying components of wages to raising inequality are identified relying on the higher-dimensional linear panel model of Abowd et al. [1999] (we will alternatively refer to this method as *two-way fixed effects model* and AKM regression). Moreover, in order to make inter-temporal comparisons, we

[22]In interpreting the graph, notice that the *absolute* predictive performance of a model has to be evaluated with respect to the unconditional wage variance. The right panel of Figure 2 is useful in order to compare the *relative* predictive performance of a model with respect to the others, but not the *absolute* one, which indeed tends to grow over time for all specifications.

adopt the same strategy of Card et al. [2013], dividing the years under study into different sub-periods. In this section, we briefly review the chosen econometric methodology, we explain the required assumptions, and we discuss the interpretation of the model.

Let i index a specific worker, t the time period, and $j = \iota(i,t)$ the firm in which i is working at t. Moreover, let y_i represent a $T \times 1$ vector of log wages, x_i a $T \times P$ matrix of time- and firm-varying individual characteristics. Then, the two-way fixed effects model can be specified as follows

$$y_{it} = x_{it}\beta + \phi_j + \eta_i + e_{it}$$

where y_{it} and x_{it} are rows of y_i and x_i, β is a $P \times 1$ vector of parameters, while ϕ_j and η_i are respectively firm-constant and time-constant components of individual wages, which are allowed to be arbitrarily correlated with any of the characteristics in x_i, and which could be not perfectly observable. We will often refer to η_i with the term *unobserved individual heterogeneity*, and to ϕ_j with *firm wage premium* or *firm wage policy*.

In the above equation e_{it} is the error term, which we assume to have an expected value equal to zero in all periods. Moreover, it is an idiosyncratic shock, which is not allowed to be correlated with any of the elements in x_i, ϕ_j and η_i. This assumption, which we define as *strict exogeneity*, can be stated formally as

$$E\left[e_{it} | x_{is}, \phi_{j=\iota(i,s)}, \eta_i\right] = 0 \quad \forall\ s,\ t$$

The above assumption rules out any pattern of endogenous mobility of workers between firms. Any realization of $\iota(i,s) = j$ should be uncorrelated with $e_{i,t}$, so that, for example, negative idiosyncratic shocks in wages should not lead to mobility towards a certain type of firms. However, any correlation between $\iota(.)$ and η_i or ϕ_j is possible, so that workers of a given type can move toward firms with certain wage policies and vice-versa. If strict exogeneity holds, the model can be consistently estimated by OLS, via inclusion of dummies for individuals' and firms' effects.

The hypothesis of exogenous worker mobility across firms, conditional on individual observable and time-constant unobservable characteristics, has been criticised, e.g. by Eeckhout and Kircher [2011] who point out that many search and matching models of the labour market are inconsistent with the additive linearity of the AKM approach. Card et al. [2013] develop several tests to support the validity of the strict exogeneity assumption. These tests have been conducted on German data (Card et al. [2013]), Portuguese data (Card et al. [2016]) and, in particular, on Italian Social Security earnings data covering a sample of firms above 50 employees over a period similar to the one analysed here (Flabbi et al. [2014]). All papers find no evidence in support of the endogenous workers' mobility hypothesis and conclude that the AKM model provides a good approximation of the wage process.

The baseline control variables included in the AKM model are a quadratic for age, a dummy for part-time workers, three dummies for occupation, the log of the number of employees, six dummies for the first five years of tenure,[23] and a full set of time fixed effects. Moreover, in order to account better for the seniority profile of earnings, we add interactions between age and occupation dummies and age and tenure dummies. To better control for business cycle volatility, we add interactions between firm size and time dummies.

Workers' fixed effects measure the personal earning capacity that is constant over time, and largely portable as individuals move to other firms during their labour market career. Instead, firm fixed effects measure how much differences in wages paid by observationally similar employers matter, keeping constant employee time-constant characteristics and other observable factors.[24] Unlike a simple average of the workers' wages in the firm, ϕ_j can be interpreted as a firm-specific wage policy because the AKM model controls for worker observed and unobserved heterogeneity, and hence accounts for the potential non-random sorting of workers to firms.[25]

[23]This discrete specification for the variable tenure is motivated by the fact that the information on tenure is censored to 1975.

[24]Notice that, as the AKM dependent variable is log wages, ϕ_j represent a *proportional* firm-specific wage premium paid by firm j to all its employees.

[25]*Firms wage premiums* can not be interpreted as indexes of firms' efficiency, since such variable depends not only on workers' skills, but also on the technology, which is endogenous to the wage (see for example Eeckhout and Kircher [2011]). Nevertheless, since the focus of this analysis is on the determinants of *wage dispersion*, rather than on firms' performance variability, the parameter ϕ_j is still highly informative for our purposes.

There are several reasons why similar firms may adopt differentiated wage policies. As highlighted by a vast stream of literature, firms might offer wages higher than the equilibrium level as part of an *exchange of gifts* with their employees (as in the *efficiency wage theory* set forth by Akerlof [1982]). Moreover, similar firms might adopt a so-called *wage posting behaviour*, offering higher wages in order to reduce the cost of vacancies (e.g. Burdett and Mortensen [1998]). Finally, firms might differ in the degree of rent-sharing, a phenomenon which Card et al. [2014] found to be small, but significant in magnitude, in the labour market analysed here.[26]

In the AKM regression each firm wage effect is computed with respect to an arbitrary reference category and, as shown by Abowd et al. [2002], it is identified only by workers who changed at least one employer within a given *mobility group*.[27] Therefore, in our analysis we have considered only the largest connected set of establishments, which usually contains, depending on the years considered, more than 95% of the sample observations. Moreover, since the estimates of firms wage premiums could be biased whenever the number of mobility episodes is low and the entire workforce is not observable, we report the main results including only for firms located in Veneto.[28]

Given the linearity of our panel model, and under the assumption of *strict exogeneity*, the total variance of log wages can be decomposed as follows

$$\begin{aligned}\mathrm{Var}(y_{it}) = {} & \mathrm{Var}(\phi_{j=\iota(i,t)}) + \mathrm{Var}(\eta_i) + \mathrm{Var}(x_{it}\beta) + \mathrm{Var}(\epsilon_{it}) + \\ & + 2\mathrm{Cov}(\phi_{j=\iota(i,t)}, x_{it}\beta + \eta_i) + 2\mathrm{Cov}(\eta_i, x_{it}\beta) \qquad (1)\end{aligned}$$

Each component in the right-hand side of the above equation can be recovered from the estimated parameters of our regression model. It follows that we can measure which are, among firm-specific, time-constant and time-varying factors, the main drivers of wage

[26] Using the years from 1995 to 2001 of the database studied here, Card et al. [2014] find an elasticity of wages with respect to profits of around 5%.

[27] By *mobility group*, or *connected set*, we intend the group containing all workers who ever worked for any of the firms in the group, and all the firms at which any of the workers in the group were ever employed.

[28] Firms excluded from the results are included in the regression, since otherwise we would have a loss in efficiency. However, we do not report results for firms outside Veneto since for such establishments we do not have information on the entire workforce. See Andrews et al. [2008] for a discussion of the effects of *limited mobility bias* in the estimates of firms-wage premiums.

dispersion, and which forces lessen their magnitude over time. With the exception of the error term, the effect of each component on the total variance is mediated by the covariance terms. Of particular interest is the covariance associated to firms' pay premiums, since it measures positive or negative sorting of individuals with given earning ability into types of firms adopting specific wage policies. Instead, the term $\text{Cov}(\eta_i, x_{it}\beta)$ measures whether workers with higher wage components related to their time-varying characteristics also tend to exhibit higher (positive covariance) or lower (negative covariance) components related to their time constant unobserved heterogeneity.

In practice it is often difficult to provide an economic intuition for which human capital factors are absorbed by unobserved heterogeneity, and what drives the sorting between time-varying and time-constant characteristics of workers, since to some extent $\text{Cov}(\eta_i, x_{it}\beta)$ is also determined by how well given workers' skills are measured by the time-varying characteristics included in the regression. Therefore, in presenting our results we more often rely on the following, more parsimonious decomposition

$$\text{Var}(y_{it}) = \text{Var}(\phi_{j=\iota(i,t)}) + \text{Var}(\eta_i + x_{it}\beta) + \text{Var}(\epsilon_{it}) + 2\text{Cov}(\phi_{j=\iota(i,t)}, x_{it}\beta + \eta_i) \qquad (2)$$

This decomposition is equivalent to the previous one, with the only exception that in equation (2) the term $\text{Var}(\eta_i + x_{it}\beta)$ captures the joint effect of workers' time-constant and time-varying characteristics. Considering only the variability of the term $(\eta_i + x_{it}\beta)$, which is a more comprehensive measure of employees' earning abilities, ease the interpretation of the results by providing a more concise information. In the analysis below we often refer to this term with the expressions *workers' portable pay component* or *workers' wage premium.*

5 Main Results

This section discusses the main results of the paper. We have estimated the two-way fixed effect model focusing on the evolution of its parameters over time, and the results of this analysis are presented in the next paragraph. Section 5.2 discusses some institutional

Table 2: **Variance Decomposition of Log Daily Wages**

Period	Var(ϕ_j)	Var(η_i)	Var($x_{it}\beta$)	Var(ϵ_{it})	2Cov(ϕ_j, $x_{it}\beta+\eta_i$)	2Cov(η_i, $x_{it}\beta$)	**TOTAL VAR.**
1982-1987	0.031	0.050	0.006	0.008	-0.019	0.007	0.083
1984-1989	0.027	0.100	0.089	0.008	-0.011	-0.121	0.092
1988-1993	0.026	0.101	0.094	0.008	0.001	-0.116	0.113
1992-1997	0.028	0.076	0.007	0.007	0.000	0.006	0.123
1996-2001	0.024	0.130	0.153	0.007	0.012	-0.195	0.131

The estimation sample is composed of all workers in the largest connected set, provided they where employed for at least four months. Results are computed only for firms located in Veneto.

features that have most likely influenced the evolution of the wage structure in the Italian case, and compares the results of the AKM regressions with those obtained on German data by Card et al. [2013]. Finally, in Section 5.3 we conduct an empirical analysis to assess the extent to which the variability of unconditional wages and of worker's pay premiums, as defined by equation (2), has been influenced by collective bargaining institutions.

5.1 Variance Decomposition from the AKM Regressions

We have calculated the variance decomposition of equation (1) on five, partially overlapping, six-years panels. In each panel, we have computed two-way fixed effects regressions controlling for human capital and aggregate shocks in wages. The coefficients associated to the regressors included in x_{it} were all significant and had the expected sign.[29] For each period, Table 2 reports the wage variance decomposition.

During the overall period considered, the total wage variance, as computed on each six-years panel, has increased from 0.83 to 0.131, growing by almost 45%.[30] When looking at the behaviour of the various components of pay, some noticeable features emerge. First, in each period, the largest contribution to the total variance derives from the joint effect of worker heterogeneity, both observed and unobserved. In general, the variance of η_i

[29] See Section 3.1 for a complete list of the regressors. Notice that the MSE for each estimation of this model is reported in Table 2. The overall R squared ranges between 0.69 and 0.77 in the male sample.

[30] To put our results in perspective, notice that Card et al. [2013] find that the total variance of male wages had risen in West Germany from 0.136 during 1985-1991, to 0.187 during 1996-2002, which translates into a 31.5% increase. Instead, in our data the male wage variance grows from 0.092 during 1984-1989 to 0.131 during 1996-2001, which translates into a 35% increase. Finally, Iranzo et al. [2008], analysing a sample of large Italian firms during the entire period 1981-1997, find the total male wage variance in the largest connected set to be 0.11.

Table 3: **Decomposition of the Total Wage Variance Evolution**

Period	Var(ϕ_j)	Var($\eta_i + x_{it}\beta$)	Var(e_{it})	2Cov($\phi_j, \eta_i + x_{it}\beta$)	**TOTAL VAR.**
1982-1987	0.031	0.063	0.008	-0.019	0.083
% of Total	*37.3*	*75.9*	*9.6*	*-22.9*	*100*
1996-2001	0.024	0.088	0.007	0.012	0.131
% of Total	*18.3*	*67.2*	*5.3*	*9.2*	*100*
Difference	-0.007	0.025	-0.001	0.031	0.048
% Δ	-25.5	33.1	-13.3	200.0	44.9
% Δ/Δ_{TOT}	**-14.6**	**52.1**	**-2.1**	**64.6**	**100.0**

Percentage changes for a given quantity z from $t-1$ to t are computed using a reference value z_r defined as $z_r = \frac{|z_t|+|z_{t-1}|}{2}$

dominates the variance of $x_{it}\beta$. However, as mentioned in the previous section, interpreting correctly what drives the relative contributions of the two workers' components is often difficult, given that unobserved heterogeneity is estimated as a residual. Also the covariance of these two terms shows a quite erratic behaviour, with high negative values in the sub-periods where the variance of $x_{it}\beta$ is relatively larger. This is the main reason why, in the rest of the paper, we tend to focus on the more parsimonious decomposition of equation (2).

A second feature of the results is that the component related to firms' wage premiums provides a smaller contribution to overall wage dispersion than worker's heterogeneity. Importantly, employers' pay policies are more relevant in the first period of the sample (1982-1987), but lose importance thereafter. Finally, the estimated correlation between firm wage effects and worker's heterogeneity (considering both its observed and unobserved components) tends to be negative in the earliest years, but it is clearly increasing over time and positive during the last period considered. Hence, there is a significant tendency towards positive sorting of firms' wage premiums with workers' overall human capital.

To show these trends more clearly, Table 3 reports the decomposition of equation (2), computed in the first and in the last panel only. In this less detailed decomposition the wage component related to a worker's time-varying observable characteristics and the component deriving from his/her time-constant unobservable skills are jointly considered.

It emerges from the table that during both periods (1982-1987 and 1996-2001) the most important determinant of total wage dispersion is the variance of the term $(\eta_i + x_{it}\beta)$, which constitutes between two thirds and three fourth of the total pay variance.[31]

The lower part of Table 3 shows the evolution of earning dispersion from the earliest to the latest panel. For each component of the total variance, we have computed the difference across samples, the percentage change, and the contribution of this change as a percentage of the change in total wage variance. Between these two periods, total wage variance has risen by almost 45%. More than 52% of this growth is driven by higher dispersion in our comprehensive measure of workers' skills. On the contrary, the dispersion in firms' wage premiums declines between the first and the last panel, providing a *negative* contribution of about 15% to the growth in wage dispersion. Finally, increasing assortative matching between highly paid workers and better paying firms provides another positive contribution to the growth in inequalities. This component represents around 64% of the total trend, even if the correlation between individual skills and ϕ_j is relatively small and close to zero in all sub-periods.

The determinants of raising assortativeness are complex and can not be fully explored within the scope of the present paper. However, at least two tendencies can be associated with this outcome.[32] First, it is tempting to relate the growth in sorting to some evolutions occurred in the Italian labour market and in its legislation since the 1980s. In common with other EU countries, Italy has indeed experienced a general trend of labour market liberalization that may have gradually reduced search and matching frictions, eventually improving allocative efficiency. For instance, in the 1980s, hiring typically involved only open-end contracts, while temporary contracts were gradually liberalized only starting from the second half of the 1990s. Moreover, during the first years of study, by the law manual workers had to be selected almost exclusively from the unemployment workers'

[31]Iranzo et al. [2008] apply AKM regression models over similar data, but with a focus on the entire period 1981-1997 (i.e., without focusing on the temporal evolution of wage inequality, and its components). Using a sample covering only firms with 50 or more employees they also find that roughly two thirds of the total wage variance is explained by worker-specific pay premiums.

[32]Changes in the composition of firms across periods represent a third possible determinant for the growth in sorting. However, Table 3.1 shows that the characteristics of firms and the employment composition across industries are quite stable over time.

lists held by the public employment service, and not via direct selection mechanisms.[33] The second possible reason behind the growth in assortativeness is more technical and linked to the strong wage compression characterizing the early 1980s. Since pay differences between skill groups were generally small in that period, the workers' portable component of wages may reflect workers' productivity less accurately during the first years of study. Thus, increased positive sorting of better paid workers to high paying firms may also be induced by a stronger relationship between wages and actual productivity across time.

In the next section we turn the discussion on how changes in the industrial relation system might have had a more direct bearing on the other two main findings of the paper, *i.e.* declining dispersion in firm wage policies and positive contribution of worker-specific wage components to the overall inequality growth. In doing so, it is useful to assess the experience of the second largest manufacturing economy in Europe (Italy) in light of what has already been documented for its manufacturing leader (Germany).

5.2 Wage Inequality and Institutions: A Comparative Perspective

Since we have used a sampling strategy and a method similar to the one that has been applied by Card et al. [2013] on German data, it is particularly interesting to compare their evidence with that provided in our study. Table 4 reports the decomposition of equation (2) applied on the results of Card et al. [2013] and on our sample, considering a comparable period of time.

Male earning dispersion has increased in Italy at a similar pace than in West Germany.[34] However, the determinants of this trend are different in the two countries. Card et al. [2013] show that, considering differences between the period 1996-2002 and the period 1985-1991, only 34% of the total growth in wage variance can be attributed to greater individual heterogeneity dispersion, while the same amount is more than 51% in the case of Veneto. Between the same periods firms' pay premiums dispersion rose by 26% in

[33] In Italy the hiring process was fully liberalized only in the early 1990s.

[34] In absolute terms, the variance of wages is higher in West Germany. Beside differences in the definition of wages across samples, another reason for this discrepancy is the sample composition. Our analysis is based on a database covering the Veneto region only, which is a more homogeneous population with a smaller and less developed tertiary sector with respect to Germany.

Table 4: **Wage Variance Evolution in Germany and Italy**

Veneto Working Histories Data, Male Sample					
Period	Var(ϕ_j)	Var($\eta_i + x_{it}\beta$)	Var(e_{it})	2Cov($\phi_j, \eta_i + x_{it}\beta$)	**TOTAL VAR.**
1984-1989	0.027	0.068	0.008	-0.011	0.092
% of Total	*29.3%*	*73.9%*	*8.7%*	*-12.0%*	*100.0%*
1996-2001	0.024	0.088	0.007	0.012	0.131
% of Total	*18.3%*	*67.2%*	*5.3%*	*9.2%*	*100.0%*
% Δ	-11.8%	25.6%	-13.3%	200.0%	35.0%
% Δ/Δ_{TOT}	**-7.7%**	**51.3%**	**-2.6%**	**59.0%**	**100.0%**
German IAB Data, Male Sample (from Card et al. [2013])					
Period	Var(ϕ_j)	Var($\eta_i + x_{it}\beta$)	Var(e_{it})	2Cov($\phi_j, \eta_i + x_{it}\beta$)	**TOTAL VAR.**
1985-1991	0.025	0.095	0.014	0.005	0.139
% of Total	*18.1%*	*67.9%*	*10.2%*	*3.8%*	*100.0%*
1996-2002	0.038	0.112	0.017	0.023	0.190
% of Total	*19.9%%*	*59.0%*	*8.9%*	*12.3%*	*100.0%*
% Δ	39.3%	16.6%	17.6%	125.5%	30.5%
% Δ/Δ_{TOT}	**24.6%**	**34.2%**	**5.5%**	**35.7%**	**100.0%**

Percentage changes for a given quantity z from $t-1$ to t are computed using a reference value z_r defined as $z_r = \frac{|z_t|+|z_{t-1}|}{2}$

Germany, while it has reduced by almost 8% in our sample. Finally, Card et al. [2013] also find that the sorting between firm-specific and employee-specific pay premiums contributed for another 36% to the overall growth in earnings inequality, which is a weaker figure than what we have documented for Italy.

Card et al. [2013] interpret their findings, and in particular the growth in firms' wage policies dispersion, as being driven by the major changes in the German industrial relation system that occurred in the early 1990s. As discussed by Dustmann et al. [2009], rather than in legislation reforms, such changes were laid out in contracts and mutual agreements between employer associations, trade unions and works councils. In response to the challenges of the post-reunification period (e.g., increasing threats of firms' off-shoring and massive migration flows), these actors allowed for an unprecedented decentralization of the German wage-setting process since the early 1990s. Deviations from industry-wide agreements through "opting-out", "opening" or "hardship" clauses were all increasingly used, even though the dominating system of industry-wide bargaining basically remained unchanged.[35] On this respect, Card et al. [2013] observe that firms' pay premiums, as computed on the 1996-2002 sample, are disproportionally lower among establishments that had *opted out* from national collective agreements, a tendency that enlarges the overall dispersion in such wage components. Thus, in Germany the growth in the variance of firm-specific wage policies (Var(ϕ_j)) was associated to a growth in the share of workers not covered by any kind of union agreement and to a rise in the number of firm-level deviations from industry-wide union agreements.

Italy's system of industrial relations shares many features of the German one, particularly for what concerns the importance of industry-wide collective bargaining. However, in many respects the Italian system has not shown the flexibility demonstrated by the German one, nor have the reforms occurred in Italy during the mid 1990s significantly weakened the influence of collective bargaining on wage setting. Italian firms have never been able to opt-out from the industry-wide settlements, adjusting wages downwardly whenever the local or firm-specific economic conditions so required (see Section 2). This may explain why, unlike in the German case, the variance of Italian firms' wage policies has not widened over time, despite the fact that also Italy has been exposed to the long-run challenges posed by the introduction of new technologies and increased international competition.[36]

Notice that, according to our estimates, the variance of firm wage policies actually decreased from the mid 1980s to the early 2000s. Unable to deviate from the industry-set minimum wages, Italian firms could still have resorted to *incremental* firm-level wage bargaining to differentiate their firm wage policies. Our data do not allow us to observe which firms or workers were covered by firm-level agreements. Nevertheless, the available evidence suggests that the incidence of firm-level agreements declined over time (e.g., Sestito and Rossi [2000]), partly as a consequence of a reduction in unionisation rates, as shown for Veneto by Vaona [2006]. The resulting standardization of compensa-

[35]See Dustmann et al. [2009] for an in-depth discussion of these and related changes occurred in the German system of industrial relations.

[36]Facilitated by the close to border location, and partly in response to competitive pressures, many Veneto firms moved their production abroad during the period considered. For example, the Romanian province of Timisoara is often called the "newest" Venetian Province, due to the large number of Veneto establishment that have opened there.

tion schemes across employers is consistent with our finding of a decreasing dispersion in firms' pay policies.

The gradual dismantling of several egalitarian institutions and practices, which took place during the 1980s, could be an additional channel explaining the higher variability in firm wage policies in the first years of our sample period, and its subsequent decline. During the 1970s, the excessive wage compression operated by the indexation mechanism of the *Scala Mobile* was distorting the wage structure defined by sectoral collective agreements.[37] Firms that wanted to provide a more adequate remuneration to the skills of their workers, most often at the team or group level, had to adopt their own firm-specific wage policy. As a result, starting from 1987, the renewals of sectoral collective agreements gradually broadened the gaps in the pay scale of the various *livelli di inquadramento*.[38] Since this evolution has probably reduced the difference between statutory minimum wages set by nation-wide contracts and the appropriate wage adjustments, it could have also contributed to reduce the adoption of differentiated firm wage policies.

Table 4 shows that the dispersion of observed and unobserved individual heterogeneity has instead been a major contributing factor to the overall wage inequality growth in the Italian case. While in principle this trend may reflect the underlying labour market forces, e.g. demand and supply of skills, in the following section we argue that such market forces have been largely "channelled" into the tracks set by the Italian system of industrial relations, particularly through the sectoral-level bargaining process. We do so by showing that the growth in individual heterogeneity dispersion has been almost entirely driven by broadened differences in pay between the job title categories (*livelli di inquadramento*) defined by industry-wide contracts.

5.3 The Impact of Collective Bargaining on Wage and Human Capital Dispersion

In this section we show that overall pay dispersion is mostly determined by *between job titles* earning variability and we link this outcome to the evolutions occurred within

[37] On this respect, the 1980s were indeed characterized by strong political pressures of white collars and intermediate managers against the excessively egalitarian policies followed by the unions (on this topic, see Manacorda [2004] among others).

[38] See CESOS [1989].

collective bargaining agreements. This result is further analysed in the Appendix A, where we show that alternative explanations for this trend have limited ground.

We have applied a variance decomposition methodology that divides total variation of a given quantity, which is partitioned into groups, into an *ecological component* (differences between groups) and an individual component (differences between members of the same group). Keeping fixed a given period t, let y_{ij} represent wages (or another quantity of interest) of worker i in group j, let n be the total number of workers, let J be the number of groups, and let n_j be the set of employees in group j. Define $\bar{y}_j$ as the average level of wages within group j, and define the within group variance as

$$\mathrm{V}_j = (\|n_j\| - 1)^{-1} \sum_{i \in n_j} (y_{it} - \bar{y}_j)^2$$

where we indicate by $\|n_j\|$ the cardinality of the set n_j (*i.e.* the number of employees in group j). Using the above notation, we can decompose the total wage variance into a *within group component*, and a *between group component* as follows

$$\mathrm{Var}(y) = \frac{1}{n-1} \left(\underbrace{\sum_{j=1}^{J} (\|n_j\| - 1)\mathrm{V}_j}_{\text{within component}} + \underbrace{\sum_{j=1}^{J} \|n_j\| \, (\bar{y}_j - \bar{y})^2}_{\text{between component}} \right) \tag{3}$$

Since the term $(\eta_i + x_{it}\beta)$ in the AKM regression model, which represents individual-specific productive abilities, is one of the main determinants of inequality, we have applied the decomposition technique defined above on this worker's *portable* wage component, albeit for comparison we have applied the same procedure also on raw earnings. In this section we present results obtained by using *livelli di inquadramento*, as defined by collective bargaining institutions, to partition the population, while the Appendix A presents results obtained by applying the same decomposition on firms.

The allocation of workers to a given *livello* is typically related to their schooling levels and to other time-invariant personal characteristics, captured by the fixed effect embedded in the worker's portable pay component. The effect of promotion to higher ladders of the scale, as well as the (fairly automatic) seniority wage premiums stipulated at each ladder by the relevant collective contract, are reflected in the time-varying component of the estimated worker premium.

Notice that individual firms can affect pay differentials between *livelli di inquadramento* only for what concerns the part above the statutory minimum wages, which are set at the industry-wide level. Moreover, by the law, employers are not allowed to downgrade workers into less remunerative job titles, an element which provides further rigidity in firms' wage adjustment decisions.[39] Given this institutional context, the *between* job titles variance can be considered an informative parameter to quantify the impact of collective bargaining on wages. A different measure is proposed by Torres et al. [2013], who include occupation dummies in an AKM model to study the effect of job title membership on wages. However, in the Italian context, our approach is more suitable for studying the influence of collective bargaining on *wage dynamics*. Indeed, in Italy the rules for assigning each worker to a job title are set by the relevant collective contract and change frequently over time. Thus, it is important to take into account the effects of such institutionally-driven shifts in the segregation of workers across various minimum wage levels,[40] something that the proposed variance decomposition, computed on a yearly basis, allows to capture to a full extent.

Before presenting the decomposition results, we provide further information on how *livelli di inquadramento* have been identified in the data. Several economic activities, despite being similar in their nature, can be regulated by more than one collective contract and the number of such industry-wide agreements, as well as the number of job titles defined by them, changes frequently over time.[41] Therefore, we have not constructed a homogeneous classification across years. We have instead considered the year-specific definition

[39] Nevertheless, employers can obviously affect the overall composition of job titles through their hiring policies, or by moving upward their current employees. We propose a test for assessing the importance of these tendencies in Appendix A.

[40] For example, several managerial occupations have started to be regulated by autonomous industry-wide collective contracts since the end of the 1980s. The resulting shift in the segregation of workers across the occupational categories and minimum wage levels defined by collective contracts is a source of challenges on how to compare and interpret the variance of job title fixed effects across time.

[41] To give an example, there are almost 40 distinct contracts covering workers in the sea transport industry, while there are only five contracts for workers in the metal-manufacturing sector.

Figure 3: **Unconditional Wage Variance and Workers' Wage Premiums Variance Decomposition Within- and Between Job Titles**

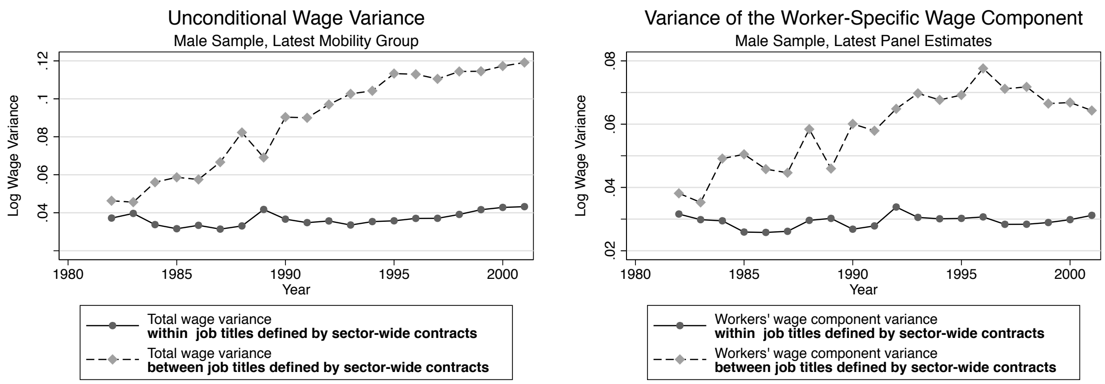

Job titles (*inquadramento* levels) are occupational categories defined within each sector-wide collective contract. In each year, we have selected only job titles represented by at least 150 workers in the largest connected set of Veneto firms, including a total number of distinct job titles between 432 (in 2001) and 520 (in 1984).

Figure 4: **Workers' Wage Premiums Variance Within- and Between Job Titles by Sector and Occupation**

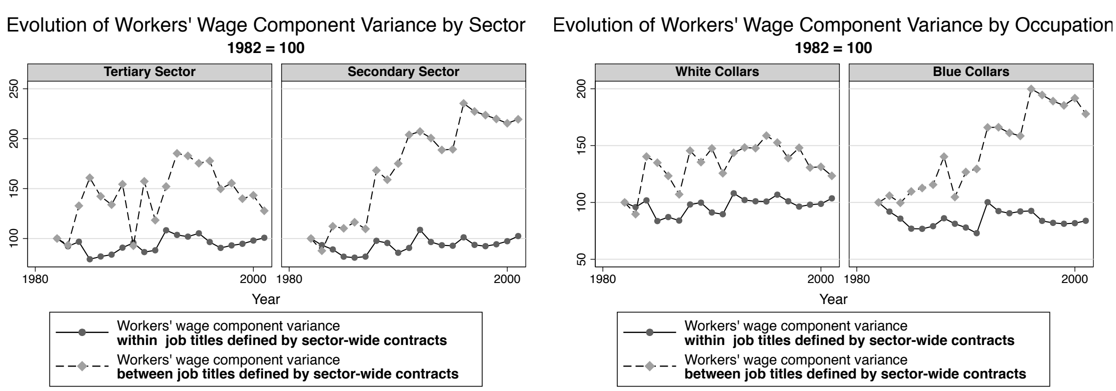

of *livelli di inquadramento*, based on their classification code. As an inclusion rule, we have adopted the criteria of considering as a legitimate job title only those for which at least 150 observations were present, in a given year, in the largest connected set of Veneto firms.[42] The total number of *livelli di inquadramento* included in the decompositions ranges between 432 (in 2001) and 520 (in 1984). Moreover, the percentage of observations which we have been able to include in our decompositions, ranges between 83% of the total in 2001 and 70% in 1986.

[42]This inclusion rule has been chosen to mitigate measurement error issues which are embedded in job titles' classification codes. When computing the variance decomposition using different thresholds, we did not find great sensitivity in the results.

Figure 3 reports the results of the variance decomposition into a between- and a within-job titles components applied year-by-year. The graphs show that practically all of the growth in the dispersion of wages and of workers' *portable* wage components is accounted for by *increased variability between livelli di inquadramento*. Indeed, both in the case of unconditional wages and of individual heterogeneity, the *between* part of the total variance shows a growing trend, with the partial exception of the second half of the 1990s, while the within component is persistently flat.[43] As a consequence, in relative terms this latter source of variation looses importance as a determinant of overall inequality. Figure 4 reports the evolution of between- and within-job titles workers' wage premiums dispersion by sector (secondary and tertiary) and by broad occupation (white and blue collars), computed by normalizing the 1982 levels of dispersion to 100.[44] A trend similar to the one implied by the right panel of Figure 3 is observed for all categories of workers, but the growth of between job titles dispersion in human capital is somewhat stronger among production workers and in the secondary sector.

A potential explanation for the trend toward higher between-job titles differences in wages could be that firms have increasingly assigned employees to higher *inquadramento levels*, as a way to raise the base wage of highly-skilled and performing workers. In the Appendix A we show that the role of this re-assignment has been relatively limited, as it can explain only a modest proportion of the between job-title variance displayed in Figure 3.[45]

The results presented in this section, together with the supplementary analysis of Appendix A, show that almost all of the inequality growth has arisen from differences in pay between job titles that are defined and protected by the industry-wide collective agreements. It remains unclear to which extent institutions have simply reacted to market forces, or whether they have represented a distortion to the wage structure. Nevertheless, we can conclude that the growth in Italian wage inequality has been allowed by

[43]An exception to this general trend is the year 1989, during which there is a drop in the between component of the total variance. However, this outcome is probably due to the measurement error induced by the mentioned change in the contracts classification code, which occurred that year.

[44]By secondary sector, we define manufacturing and constructions sector. The primary sector (agriculture, forestry, fishing and mining) is excluded from these computations. The service sector is defined as the residual category.

[45]Notice that, as mentioned, this re-assignment can only be operated upwardly, i.e. it is not possible to downgrade a worker to a lower *livello di inquadramento*.

the opening of the pay gaps between the various *livelli di inquadramento* stipulated in a fairly centralized way at each industry-wide contract renewal, combined with the gradual dismantling of the egalitarian wage indexation system since the mid 1980s.

6 Conclusions

In this paper we have analysed the evolution of Italian wage inequality over a two decades period, documenting a substantial growth in several measures of pay dispersion. To interpret this trend, we have decomposed the wage variance into components capturing heterogeneity in firm pay policies, heterogeneity in workers' time-varying and time-constant characteristics, as well as their sorting. We have found that earnings dispersion has been mostly driven by differences in the workers' *portable* component of wages. Instead, the variability of employer-specific pay premiums has reduced over time.

Our results are different from evidences documented for other countries, and Germany in particular. On this respect, we have provided an indirect support to the conclusions of Card et al. [2013]. These authors report evidence of a growth in firms' pay premiums dispersion. They attribute this finding to firm-level deviations from the dispositions of industry-wide collective agreements (e.g., the *opting-out* clauses), which were allowed by the German system and became increasingly used since the mid 1990s. We have documented the lack of such a flexible adaptation process in a similar manufacturing-oriented economy, which has undergone qualitatively different reforms in its system of industrial relations. Italian firms have been unable to apply heterogeneous pay policies, and to circumvent the constraints to wage dynamics imposed by the sectoral level of bargaining.

To shed further light on the role played by collective bargaining in the observed inequality trend, we have analysed the evolution of pay differentials across so-called *livelli di inquadramento*. These are job titles defined by nation-wide sectoral collective agreements, for which specific minimum wages hold and apply regardless of a worker's union membership. A simple variance-decomposition exercise allowed us to show that the growth in both, wage and human capital dispersion, has almost entirely occurred between such job titles.

Overall, our results show that market forces have been largely "channelled" into the tight tracks set by the rules governing the country's fairly centralized system of industrial relations. Moreover, during the overall period considered, firms have been granted very limited margins of wage flexibility. The extent to which the bargaining system may have been able to provide the adequate signals about the *appropriate* wage adjustments, i.e. the adjustment required by underlying market forces, remains an open question.

Appendix

A Related Evidence on Between Job Titles Pay Dispersion

In this section we provide further evidences related to the evolution of wage inequality between the so-called *livelli di inquadramento*. In particular, we test whether this type of dispersion has been driven by a trend in the composition of the labour force, which could have become more likely to be employed at relatively low-paid (or high paid) occupations. Moreover, we test whether differences in wages and human capital arise mostly within or between firms.

The left panel of Figure A.1 shows the proportion of workers within each quartile of the job titles' average pay distribution. In constructing the graph, we have computed year-by-year the average wage within each job title, separately considering workers in the secondary and tertiary sectors. For each of these two sectors, we have classified each job title according to the quartile of the job titles' average pay distribution to which it belongs.[46] Then, we have computed year-by-year the proportion of workers within each quartile group of job titles.

In the right panel of Figure A.1 we replicate the analysis described above, this time considering the worker-specific component of the wage, as estimated by the AKM regression model. In particular, we have computed the average level of *skills* (*i.e.* observed and unobserved individual heterogeneity) within each job title. We have used this information to rank job titles into quartiles, and we have computed the proportion of workers within

[46] The job titles' average pay distribution is not weighted by the number of observations within each job title category. Thus a given percentile of the job titles' average pay distribution can be quite different from the same percentile of the wage distribution.

Figure A.1: **Proportion of Workers Within Quartiles of the Job Titles' Average Pay and Skills Distribution**

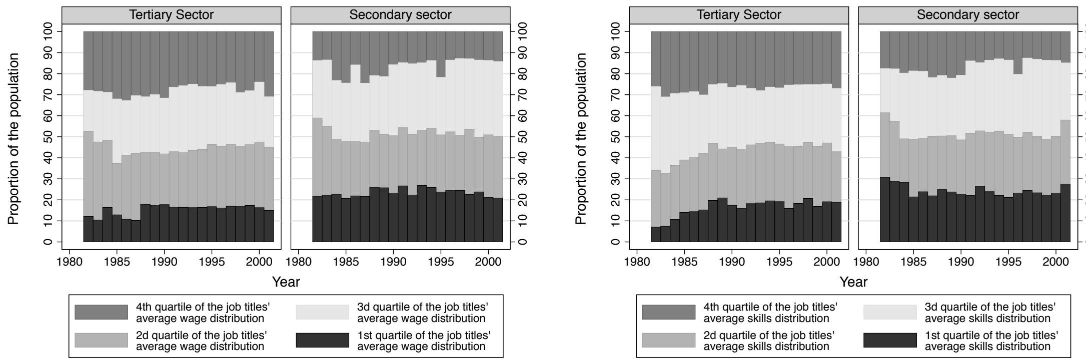

The 1st quartile refers to the bottom 25th percentiles of the job titles' average pay (or, in the left panel, *workers' premiums*) distribution. Similarly, the 4th quartile refers to job titles characterized by an average wage (human capital) higher than the 75th percentile of the job titles' average pay (human capital) distribution.

each quartile group.

The left panel of Figure A.1 shows that the proportion of workers within each quartile of the job title pay distribution has been fairly constant during the overall period considered. There are some exception to this general trend. In particular, in the years 1982 and 1983 the proportion of secondary sector workers belonging to the upper quartiles of the job title distribution is quite low. This tendency is probably induced by the wage adjustment mechanisms in place up until 1984, which were shifting upward wages at the bottom of the pay scale. Moreover, during the same two years the proportion of workers in the two lowest quartiles of the tertiary sector was smaller than in subsequent years. During the years between 1984 and the beginning of the 1990s, both in the secondary and tertiary sectors there is a small growth in the proportion of workers in the highest quartile of job titles. However, at least in the secondary sector, this tendency is only cyclical, given that, from the early 1990s onwards, this same proportion decreases to levels similar to the ones in place during the early 1980s. Finally, in the service sector only, there is a relatively persistent growth of the proportion of workers belonging to the lowest quartile of the job title distribution. However, this growth is quite small in magnitude. Other discrepancies across time tend to be year-specific, and are most likely attributable to differences in the job title classification codes from one year to the other.

The right panel of Figure A.1 shows that even when differences in the characteristics of employers are accounted for, most tendencies remain similar to the ones described in the left panel. The growth in the proportion of service sector workers belonging to the first quartile of job titles is confirmed and it is even stronger in this case. However, also here other trends seem to be either year specific, or relatively small in magnitude.

Overall, by analysing the composition of job title categories across time, we can conclude that the main channel driving greater wage dispersion is linked to differences in how the same occupations are rewarded across time. Thus, with the partial exception of the service sector, where employers have become more likely to hire workers belonging to less remunerative job title categories, there is no evidence of a process of *polarization* of the workforce.

The growth of pay differentials between job titles, which we have documented in Section 5.3, could also derive from a process of segregation of the more qualified workers into given enclaves of firms.[47] We test this hypothesis computing year-by-year the decomposition of equation (3), this time using firms as the partitioning group of the population.

Figure A.2 reports the within- and between-firms variance decomposition, applied on wages and on the estimated individual heterogeneity of the workforce. The left panel of the figure shows that wage variation was almost equally split into a within- and a between-firms component during the early 1980s. Since then, the importance of earnings variance among co-workers rises sharply with respect to differences in average wages between plants. Nevertheless, the years under study are characterized also by a small growth in the unconditional wage variance *between firms*. This latter process can be mostly ascribed to increased sorting of workers' wage premiums with firms' pay premiums, a result which emerged from the AKM variance decomposition.

In the right panel of Figure A.2 we compute the same variance decomposition using workers' wage premiums alone, instead of total wages. It emerges that the dispersion *between*

[47] Such market-driven process would then probably be reflected in collective bargaining dynamics, given that the more skill-intensive firms could be able to grant better economic conditions to selected groups of job titles. On the other hand, if pay differentials between firms are low, despite a general growth in job title heterogeneity, we may think that firms are constrained by the sectoral bargaining standards, given that most of the inequality growth occurs within establishments, instead of across them.

Figure A.2: **Within- and Between-Firms Decomposition of Unconditional Wage Variance and Workers' Wage Premiums Variance**

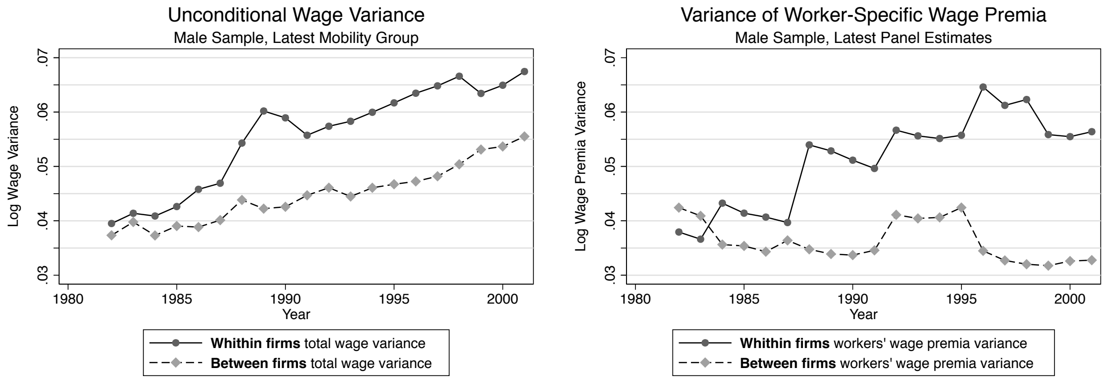

Worker's wage premiums variance is defined as $\text{Var}(\eta_i + x_{it}\beta)$. Since each panel that we have constructed is partially overlapping, for each year we report only estimates of $\text{Var}(\eta_i + x_{it}\beta)$ from the latest available period. For each year, the unconditional wage variance is computed on the largest connected set in the latest panel. Only firms located in Veneto are considered.

employers of this component of the pay, which controls for heterogeneity in firms' wage residuals, has even declined over the entire period. Therefore, we find no evidence of greater segregation of workers' skills across employers.

The low level of between-firms pay dispersion documented here is coherent with previous studies on Italy (such as Iranzo et al. [2008]), but it is a quite peculiar result when compared with evidences available for other European countries and the US.[48] Moreover, our finding is particularly robust, given that the sample includes also very small firms and the private service sector, which are two categories whose exclusion could drive the estimates of between-plants pay differentials down.[49]

Overall Italy is not characterized by strong pay differences between firms, which could have been relevant if, for example, greater dispersion in productive performance across employers, often considered an outcome of technological changes and international competition, had induced greater heterogeneity in wages between plants. Instead, the relevance of pay dispersion between job titles, which is documented by Figure 3, suggests that, in the Italian case, the growth of inequality has entirely occurred within the collective

[48]Among studies focusing on between-plant wage inequality in other countries, see for example Faggio et al. [2010] on UK. Card et al. [2013] show that in Germany firms pay premiums dispersion has risen over time also when accounting for employees' sorting across plants.

[49]Firm-size wage premiums (found by Scoppa [2014] also in Italy) may induce an underestimation of between plants wage differentials when small firms are excluded.

bargaining framework. Based on our results, we can conclude that over the years such institution has granted more heterogeneous conditions for selected categories of workers (*i.e.* job titles), while it has provided limited margins of flexibility for the firms.

B Other Figures and Tables

Figure B.1: **Evolution of GDP and Employment in the Period of Study**

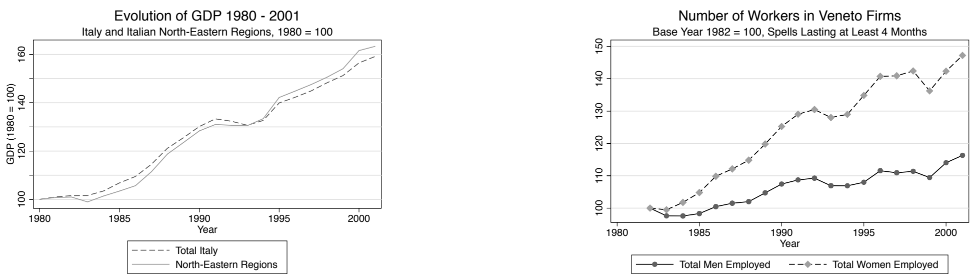

Source: ISTAT (Italian Statistical Office)

Figure B.2: **Long-Run Evolution of Gross Weekly Wages**

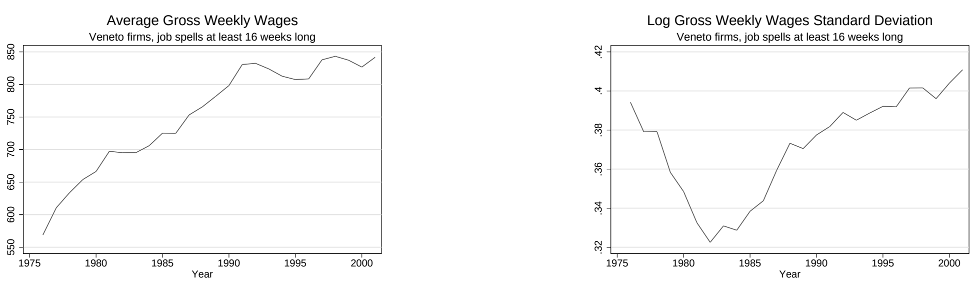

References

Abowd, J., F. Kramarz, and D. Margolis (1999): "High Wage Workers and High Wage Firms," *Econometrica*, 67, 251–333.

Abowd, J. M., R. H. Creecy, and F. Kramarz (2002): "Computing Person and Firm Effects Using Linked Longitudinal Employer-Employee Data," Center for economic studies, u.s. census bureau.

Acemoglu, D. and D. Autor (2011): *Skills, Tasks and Technologies: Implications for Employment and Earnings*, Elsevier, vol. 4 of *Handbook of Labor Economics*, chap. 12, 1043–1171.

Akerlof, G. A. (1982): "Labor Contracts as Partial Gift Exchange," *The Quarterly Journal of Economics*, 97, 543–69.

Andrews, M. J., L. Gill, T. Schank, and R. Upward (2008): "High Wage Workers and Low Wage Firms: Negative Assortative Matching or Limited Mobility Bias?" *Journal of the Royal Statistical Society Series A*, 171, 673–697.

Autor, D. H., L. F. Katz, and M. S. Kearney (2008): "Trends in U.S. Wage Inequality: Revising the Revisionists," *The Review of Economics and Statistics*, 90, 300–323.

Brandolini, A., P. Cipollone, and P. Sestito (2002): "Earnings Dispersion, Low Pay and Household Poverty in Italy, 1977-1998," in *The Economics of Rising Inequalities*, Oxford University Press.

Burdett, K. and D. T. Mortensen (1998): "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, 39, 257–73.

Cappellari, L. (2004): "The Dynamics and Inequality of Italian Men's Earnings: Long-term Changes or Transitory Fluctuations?" *Journal of Human Resources*, 39, 475–499.

Cappellari, L. and M. Leonardi (2016): "Earnings Instability and Tenure," *The Scandinavian Journal of Economics*, 118, 202–234.

Card, D., A. R. Cardoso, and P. Kline (2016): "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women," *The Quarterly Journal of Economics*, 131, 633–686.

Card, D., F. Devicienti, and A. Maida (2014): "Rent-sharing, Holdup, and Wages: Evidence from Matched Panel Data," *Review of Economic Studies*, 81, 84–111.

Card, D., J. Heining, and P. Kline (2013): "Workplace Heterogeneity and the Rise of West German Wage Inequality," *The Quarterly Journal of Economics*, 128, 967–1015.

Casadio, P. (2003): "Wage Formation in the Italian Private Sector after the 1992-93 Income Policy Agreements," in *Institutions and Wage Formation in the New Europe*, Edward Elgar.

CESOS (1989): "Le relazioni sindacali in Italia, Rapporto 1987-1988," Edizioni lavoro, CESOS.

Dell'Aringa, C. and C. Lucifora (1994): "Collective Bargaining and Relative Earnings in Italy," *European Journal of Political Economy*, 10, 727–747.

Devicienti, F. and A. Borgarello (2001): "Trends in the Italian Earnings Distribution, 1985-1996," Working paper, LABORatorio Revelli.

Devicienti, F., B. Fanfani, and A. Maida (2016): "Collective Bargaining and the Evolution of Wage Inequality in Italy," IZA Discussion Papers 10293.

Devicienti, F., A. Maida, and L. Pacelli (2008): "The Resurrection of the Italian Wage Curve," *Economics Letters*, 98, 335–341.

Di Nardo, J., N. Fortin, and T. Lemieux (1996): "Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach," *Econometrica*, 64(5), 1001–1044.

Dustmann, C., B. Fitzenberger, U. Schonberg, and A. Spitz-Oener (2014): "From Sick Man of Europe to Economic Superstar: Germany's Resurgent Economy," *Journal of Economic Perspectives*, 28, 167–88.

Dustmann, C., J. Ludsteck, and U. Schnberg (2009): "Revisiting the German Wage Structure," *The Quarterly Journal of Economics*, 124, 843–881.

Eeckhout, J. and P. Kircher (2011): "Identifying Sorting–In Theory," *Review of Economic Studies*, 78, 872–906.

Erickson, C. and A. Ichino (1995): "Wage Differentials in Italy: Market Forces, Institutions, and Inflation," in *Differences and Changes in Wage Structures*, National Bureau of Economic Research, Inc, NBER Chapters, 265–306.

Faggio, G., K. G. Salvanes, and J. V. Reenen (2010): "The Evolution of Inequality in Productivity and Wages: Panel Data Evidence," *Industrial and Corporate Change*, 19, 1919–1951.

Flabbi, L., M. Macis, A. Moro, and F. Schivardi (2014): "Do Female Executives Make a Difference? The Impact of Female Leadership on Gender Gaps and Firm Performance," IZA Discussion Papers 8602.

Iranzo, S., F. Schivardi, and E. Tosetti (2008): "Skill Dispersion and Firm Productivity: An Analysis with Employer-Employee Matched Data," *Journal of Labor Economics*, 26, 247–285.

Katz, L. F. and K. M. Murphy (1992): "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *The Quarterly Journal of Economics*, 107, 35–78.

Leonardi, M., M. Pellizzari, and D. Tabasso (2015): "Wage Compression within the Firm," IZA Discussion Papers 9254.

Manacorda, M. (2004): "Can the Scala Mobile Explain the Fall and Rise of Earnings Inequality in Italy? A Semiparametric Analysis, 1977-1993," *Journal of Labor Economics*, 22, 585–614.

Naticchioni, P. and A. Ricci (2009): "Falling Educational Wage Premia in Italy," *QA - Rivista dell'Associazione Rossi-Doria*, 4.

Piketty, T. and E. Saez (2003): "Income Inequality In The United States, 1913-1998," *The Quarterly Journal of Economics*, 118, 1–39.

Scoppa, V. (2014): "Firm Size and Wages in Italy: Evidence from Exogenous Job Displacements," *Manchester School*, 82, 677–700.

Sestito, P. and F. Rossi (2000): "Contrattazione aziendale, struttura negoziale e determinazione decentrata del salario," *Lavoro e Relazioni Industriali*, 2, 129–183.

Torres, S., P. Portugal, J. T. Addison, and P. Guimaraes (2013): "The Sources of Wage Variation: A Three-Way High-Dimensional Fixed Effects Regression Model," IZA Discussion Papers 7276.

Vaona, A. (2006): "L'evoluzione recente dei tassi di sindacalizzazione in Veneto," *Quaderni di Rassegna Sindacale*, 3.

Chapter III

The Gender Wage Gap Among Italian Employees Evidences from the ISFOL PLUS Database*

Abstract

In this paper, we measure the amount of gender discrimination in the Italian labour market, using the ISFOL PLUS database and covering the period between 2005 and 2010. We adopt a quantile regression decomposition methodology, in order to measure the gender wage gap across the entire distribution of earnings. Moreover, we test the robustness of our results by employing panel techniques. The results show that, other things being equal, female employees earn around 10% less than men employees. This percentage is increasing with wages, reaching levels higher than 20% toward the top of the earnings distribution. We show that this gap is increasing with age, and that this trend is driven by an improvement of market potential among the younger generation of women. Moreover, we show that the pay gap was reducing between 2006 and 2008, but it has increased since then. Finally, we show that most of the gender wage differences are not attributable to individual characteristics, nor to segregation of women into less remunerated occupations. Therefore, most of the gender differences in income remain unexplained by the rich set of observable characteristics included in our data, and should be attributed to other labour market dynamics that are disadvantageous for women.

JEL Codes: J00, J16, J31, J7. **Keywords**: Gender Wage Gap, Discrimination, Quantile Regression, Oaxaca Decomposition.

*I would like to thank Carolina Castagnetti, Emiliano Mandrone and Francesco Devicienti for their support and the useful comments. A version of this paper has been published under the same title in the peer-reviewd series *ISFOL Research Papers* (No.26, 2015).

1 Introduction

Nowadays most of the Western Countries, and Italy in particular, are facing the strong challenges posed by demographic trends. Low fertility rates, as well as increased life expectancy rates, are putting pressure on the sustainability of welfare systems. In order to respond to increasing levels of economic dependency,[1] the most desirable paths are productivity improvements and a growth in labour force participation, since both solutions are needed to maintain the current level of welfare benefits.

In this context, it is important to investigate whether discrimination against groups of the population is playing an important role in the labour market. Discrimination represents an inefficiency, and it reduces productivity by determining a sub-optimal allocation of resources. Moreover, discrimination reduces labour force participation, since it lowers the incentives to work for those who suffer from it. In particular, discrimination against women is very harmful, since it involves a large proportion of the labour force.

In this paper, we are going to measure the amount of gender discrimination among Italian employees, using the 2005 to 2010 waves of the ISFOL PLUS database. For this

[1] Economic dependency is defined as the ratio between the working population, which is financing welfare systems, and the inactive population.

purpose, we are going to employ a quantile regression based decomposition, which is a methodology that allows to measure differences in discrimination among higher and lower remunerated job positions. We will review and derive carefully the chosen estimator. Moreover, we are going to discuss in detail which is the definition of discrimination adopted, what are the assumptions needed, and whether the findings are robust when imposing less demanding assumptions. Finally, we will try to derive some conclusions on what is driving the gender wage gap in the Italian labour market.

The estimated level of the gender pay gap in Italy amounts to around 10%, except at the upper tail of the wage distribution, where it increases substantially. Overall, our results are quite similar to the ones found by Christofides, L., A. Polycarpou, and K. Vrachimis (2010) and by Di Tommaso and Piazzalunga (2013). Both studies make use of the quantile regression based decomposition, analyse years that are covered by our data, and apply the Heckman (1979) procedure to take into account selection problems. Instead of the Heckman correction, here we will carry out several tests on the sensitivity and robustness of our results. Taking advantage of the longitudinal structure of the data, we will apply panel techniques, and in particular the Hausmann and Taylor (1981) model, in order to take into account the problem of unobserved individual heterogeneity. However, one should be aware that correlation between *time invariant* characteristics and individual fixed effects is not allowed under the Hausman-Taylor model. In particular, we won't be able to obtain estimates of discrimination that are consistent under correlation of gender to unobservable abilities. However, we can rely on the fact that the rich set of information contained in the PLUS samples allows us to build a good model of wage prediction, and we will interpret our measures of discrimination as a composite residual effect.

We will show that the amount of gender discrimination in the Italian labour market is substantial, and explains almost all of the gender differences in earnings. We will show some interestenting patterns that may help understanding better where such gender differences originate. First, we will show that the gender wage gap is influenced by the business cycle, since it has not been constant across years. Second, we will show that it is increasing with age. Our model shows that this pattern is determined partly by a lower level of discrimination among younger workers, and in part by better market potential of younger women. However, the positive relation between seniority and discrimination could be driven by several dynamics, which we can't fully take into account. Finally, we will show that most of the gender differences in pay are determined *within* occupations and *within* sectors. That is, most of the differences in the Italian labour market are not the result of segregation of women into less remunerative sectors or occupations, but rather they are determined by different payment structures for similar jobs.

2 Theoretical Framework

In economic literature, discrimination against sub-groups of the population had been an almost neglected topic until Becker seminal work, The Economics of Discrimination, published in 1957. According to his approach, discrimination in the marketplace can be modelled as an implicit transaction cost, by introducing a so-called taste for discrimination:

Discrimination is commonly associated with disutility caused by contact with some individuals, and this interpretation is followed here. [...] To the employer [discrimination] represents a non-monetary cost of production, to the employee a non-monetary cost of employment, and to the consumer a non-monetary cost of consumption (Becker, 1973, p. 15).

An empirical methodology to quantify the amount of discrimination in the marketplace has been introduced by Oaxaca (1973) and Blinder (1973). Since then, a growing number of publications have addressed this problem, and so called decomposition methodologies have been applied in many different contexts.[2] The Oaxaca-Blinder approach is based on the following definition of discrimination against sub-groups of the population

[2] Hereafter, when using the term wage structure, we will refer to the pay schedule faced by individuals, given their set of skills. As we will see in a moment, we are assuming no general equilibrium effects.

(1) $$D = \frac{(w_m/w_f) - (\tilde{w}_m/\tilde{w}_f)^{ND}}{(\tilde{w}_m/\tilde{w}_f)^{ND}}$$

where w_i are wages observed for the (gender) group $i \in \{m, f\}$, and the superscript ND denotes the hypotetical and unkown wage ratio that would be observed in the absence discrimination. In a competitive equilibrium, due to the well known theory of cost minimization, the non-discriminatory ratio of wages $(\tilde{w}_m/\tilde{w}_f)^{ND}$ would equate the ratio of male and female marginal products. One way of approximating such marginal contributions is by estimating a mincerian wage equation, where labour income is considered a true measure of productivity, and it is predicted using a series of controls for human capital and other relevant individual attributes. Such wage equation usually takes the following semi-logarithmic functional form $\ln w_i = x_i\beta_i + \varepsilon_i$, where x_i is a $n_i \times p$-matrix containing a constant, while β_i is a p-vector of coefficients.

If there was no discrimination in the labour market, one would expect that individual characteristics were rewarded equally across gender groups. Stated differently, in a non-discriminatory market the wage structure faced by males would also apply to females.[3]

---

[3] Hereafter, when using the term *wage structure*, we will refer to the pay schedule faced by individuals, given their set of skills. As we will discuss in more detail in a moment, we are assuming *no general equilibrium effects*. Broadly speaking, we are assuming that males are rewarded according to the competitive prices for skills, while women are discriminated. Instead, one could also construct the counterfactual distribution of wages where men skills are evaluated according to the female pay schedule, or according to a weighted combination of the two pays schedules.

Then, using the properties of OLS, it is straightforward to decompose differences in income as follows

(2) $$\ln w_m - \ln w_f = \overline{x_m}\widehat{\beta_m} - \overline{x_f}\widehat{\beta_m} + \overline{x_f}\widehat{\beta_m} - \overline{x_f}\widehat{\beta_f} = \left(\overline{x_m} - \overline{x_f}\right)\widehat{\beta_m} + \overline{x_f}(\widehat{\beta_f} - \widehat{\beta_m})$$

where bars represent mean values, while parameters estimated by applying OLS separately in the male and in the female samples are denoted with a hat. $\overline{x_f}\widehat{\beta_m}$ is a counterfactual wage, which measures the average wage that women would earn, had they been paid as men are. Equation (2) is the classical Oaxaca-Blinder decomposition. The term $\left(\overline{x_m} - \overline{x_f}\right)\widehat{\beta_m}$ is the so called characteristics effect, which is coherent with the wage prediction model, since it is driven by mean differences in individual skills among the two groups. Instead, the second addend of (2) is the coefficient (or wage structure) effect. It measures differences in the way gender groups are rewarded for the same characteristics.

Using the decomposition (2), a logarithmic equivalent of the coefficient D in (1) can be defined as[4]

[4] This definition of discrimination follows from the fact that $\overline{x_f}\widehat{\beta_m}$ is choosen as the counterfactual wage. See footnote 4 and the description of the no general equilibrium effects assumption for a more detailed discussion of this point.

$$\ln(D+1) = \ln(w_m/w_f) - \ln(\widetilde{w}_m/\widetilde{w}_f)^{ND} = (\ln w_m - \ln w_f) - \left(\overline{x_m} - \overline{x_f}\right)\widehat{\beta_m}$$

$$= \overline{x_f}(\widehat{\beta_f} - \widehat{\beta_m})$$

It is important to discuss what are the relevant assumptions that have to be satisfied, in order to correctly identify this discrimination coefficient. In particular, there are three conditions that have to be imposed.

- **No general equilibrium effects** We are assuming that the counterfactual (non-discriminatory) wage structure β^* will not be affected by the removal of discrimination from the labour market. In particular, in the context of the decomposition (2), we have assumed that the male pay schedule would prevail in a fair labour market, so that $\beta^* = \beta_m$.[5]
- **Overlapping support** Either we assume that the estimated regression coefficients can be extended to combinations of covariates not observable in the data, or we need to restrict attention to combination of characteristics observable among both men and women.[6]
- **Ignorability** This is the most important assumption and, in its more general

5 See Oaxaca and Ransom (1994) for a more detailed discussion of this assumption, and for possible alternatives.

6 This problem has been addressed explicitly by Nopo (2008), who introduces a method based on a matching algorithm. He proposes a four-fold wage decomposition, where the additional components represent wage differences between matched and unmatched observations in each gender group.

formulation, can be stated as follows. Let $F(.)$ represent the conditional distribution of the error term. Then

$$(3) \qquad F(\varepsilon_m|x_m) = F\left(\varepsilon_f\middle|x_f\right) = F(\varepsilon|\mathrm{x})$$

Classical exogeneity is not required, as long as the conditional distribution of the error term is the same among men and women.[7] However, imposing exogeneity of the regressors in the wage equation of both groups is a sufficient condition for the ignorability assumption to be satisfied. Chernozhukov, Fern´andez-Val, and Melly (2013) have shown that, under ignorability, the discrimination coefficient can have a causal interpretation, in the sense that it reflects solely gender differences attributable to the wage structures.

It is important to recognize that many factors are likely to induce a violation of condition (3). On one hand, group-differences in unobservable individual heterogeneity could contribute substantially to the determination of the wage gap. To some extent, this bias can be limited through the use of panel techniques, like a within transformation of

[7] Such assumption is satisfied in the context of exougenous policies, when the members of each group we want to compare are randomly selected. Whenever this is not the case, a possible strategy is to adopt some form of correction for self-selection bias, such as the procedure developed by Heckman (1979). See Buchinsky (1998) for an analogous procedure in the context of quantile regression. See Di Tommaso and Piazzalunga (2013) for a recent estimation of the Italian gender wage gap using the the Heckman-correction.

the data. However, these methods, which will be discussed and implemented in Section 6, have further limitations. Indeed, the more robust estimators often require the use of some instrument, which might be difficult to be found.[8]

A second reason why condition (3) could be violated is due to the presence of firm-specific heterogeneity in wage compensation schemes. In their seminal work, Abowd, Kramarz, and Margolis (1999) have found evidence suggesting that the more profitable firms are those who tend to pay higher wages. In our context, a bias would arise if there were gender differences in the way workers sort into high-wage firms.

Card, Cardoso, and Kline (2013), using a Portuguese matched employer-employee database, have found that, even after controlling for firm fixed effects, the relation between firm profits and wages is weaker for women. This suggests that women do not systematically work for less profitable firm, but rather that they gain less than men from firms' profits. However, it is not clear whether this tendency could be found also in the Italian labour market, especially because female labour force participation is higher in Portugal. Using a similar matched employer-employee database for Italy, which covers

8 In particular, by using a simple a fixed effect regression, the coefficients of time invariant regressors, such as schooling, can't be recovered. A possible solution is to estimate the regression model proposed by Hausman and Taylor (1981), and to control the endogeneity of time invariant regressors using some instruments. See Polachek and Kim (1994) for a review of panel techniques for the estimation of the gender earning gap.

the period 1996-2003, Matano and Naticchioni (2013) have found no evidence of under-representation of women in more profitable firms. Both of the evidences above suggest that the role played by firm fixed-effects in the determination of the gender earnings gap might be limited.

Finally, a reason that could explain the gender pay gap, but that is not accounted for by our framework, is given by the possible presence of gender-specific differences in labor supply elasticity to the firm. Indeed, one driver of discrimination could arise in monopsonistic labour markerts,[9] if women had a rigid labour supply at the firm level. The monopsonistic employer could extract more rents from employees whith a rigid supply, and women could be disadvantaged for this reason. This type of discrimination is usually referred to as *robinsonian discrimination*,[10] and is conceptually different from the kind of discrimination defined by Becker. Indeed, according to the monopsonistic framework, discrimination would be an equilibrium outcome, rather then the product of inefficient and discriminatory markets. Sulis (2011) calculates the wage elasticities to the recruitment and separation rates for Italy, analysing the period between 1985 and 1996.

---

[9] A monopsony is a market with many sellers and only one buyer, who can extract a rent in a similar way to that of a monopolist. The monopsony could be a realistic framework for modelling labour markets, mainly due to the presence of frictions which reduce the mobility of workers and to the presence of excess labour supply. For a review of the recent letature on monopsony see Ashenfelter, Farber and Ransom (2010).

[10] See Ransom and Oaxaca (2010) for a discussion and a more precise definition of the concept of *robinsonian discrimination*.

He finds that women labour supply to the firm is significantly more rigid than the one of men. This result seems to provide indirect evidence supporting the hypothesis of *robinsonian discrimination*, even if the link between wage elasticities to the firm and discrimination, to the best of our knowledge, has not yet been established in the literature.

In general, given the nature of our data, which lack information about individual employers, we will not be able to test more elaborate hypothesis that could explain the presence of gender discrimination. Therefore, we should be aware that a decomposition of the wage structure effect into a component due to sorting of workers, a component due to labour market frictions, and a component attributable to pure bargaining effects, is not feasible. In general, the estimated discrimination coefficient should be interpreted more like a composite residual effect, rather than a simple employers' disutility parameter. Therefore, policies designed to reduce the levels of wage discrimination should take into account the complexity of this problem. Gender differences in pay likely originate from a variety of labour market characteristics, all of which should be taken into account in order to design effective policies.

3 Econometric Methodology

We are now going to illustrate the econometric methodology used to identify the discrimination coefficient described in the previous section. The chosen estimator, originally developed by Machado and Mata (2005) and Melly (2005), is usually called quantile regression decomposition. The main advantage of this technique, with respect to the traditional OLS decomposition, is that it allows us to study the wage gap along the distribution of income, and not only at the average level.

Quantile regression, as developed by Koenker and Bassett (1978), is based on the notion of conditional quantiles of the dependent variable y (log wages), given the covariates x (individual characteristics).[11] Such relation is described by the following function

$$Q_y(\theta|x) = x\beta(\theta) \tag{4}$$

where Q_y is used to indicate θth conditional quantile of y. The parameter β is estimated as the solution to the following problem

$$\widehat{\beta}(\theta) = \operatorname{argmin}_{\beta \in R^p} \sum_{k=1}^{n} \varrho_\theta (y_k - x_k\beta)$$

where n represents the sample size, the parameter space is given by R^p, and the ϱ_θ loss function takes the form

[11] Here, y is a n×1 vector and x is a n×p matrix.

$$\varrho_\theta(u) = u[\theta - I(u < 0)] \qquad\qquad \theta \in [0, 1]$$

where I(.) represents the indicator function, equal to one if the term inside the brackets is true and zero otherwise. In a finite sample, the number of distinct regression coefficients $\widehat{\beta}(\theta)$ that can be estimated is finite. Let $\Theta = \{\theta_0 = 0, \theta_1,\ldots,\theta_J = 1\}$ be the set of points where the solution changes, and notice that $\widehat{\beta}(\theta_j)$ prevails from θ_{j-1} to θ_j for j = 1,..., J. Moreover, let $\widehat{\beta} = \left[\widehat{\beta}(\theta_1), \ldots, \widehat{\beta}(\theta_j), \ldots, \widehat{\beta}(\theta_J)\right]$ be the vector of all different quantile regression coefficients. Using such solutions, we can build a model for the estimated conditional quantiles $\widehat{Q}_Y$, in order to recover a conditional distribution of income.[12] Then, following Melly (2005), we will be able to construct a non-discriminatory wage structure using a particular unconditional (marginal) distribution. Let F_Y represent the distribution function of the random variable y. The θth quantile of y is defined as $y_\theta = F^{-1}(\theta)$. Therefore

$$F_Y(y_\theta) = P(y \le y_\theta) = \int_{-\infty}^{y_\theta} dF_Y(z) = \theta$$

[12] Using the entire set of solutions increases the risk of quantile-crossing. That is, $\widehat{Q}_Y(\theta|x)$ could be non-increasing in θ when evaluated at a given x. In general, the larger the number of solutions used to approximate the conditional distribution of y, and the smaller the number of observations available, the greater becomes the risk of quantile-crossing. Notice however that using the entire set of solutions is not necessary, since a sufficiently large set $\Omega \subseteq \Theta$ of quantiles, which can be drown from a uniform on [0, 1], will produce valid results.

The above result can also be obtained by integrating θ over the interval [0, $F_Y(y_\theta)$].[13] Indeed, we have that

$$\int_{-\infty}^{F_Y(y_\theta)} d\theta = \theta$$

Consider now the following indicator function, defined as

$$I[F_Y^{-1}(\theta) \leq y_\theta] = \begin{cases} 1 & if\ \theta \in \{z \in (0,1): F_Y^{-1}(\theta) \leq y_\theta\} \\ 0 & otherwise \end{cases}$$

Notice that integrating such function, with respect to θ, over the interval [0,1], is equivalent to integrate θ over the interval [0, $F_Y(y_\theta)$]. Given the above definitions, we can conclude that

(5) $$F_Y(y_\theta) = \theta \Leftrightarrow \int_0^1 I[F_Y^{-1}(\theta) \leq y_\theta] d\theta = \theta$$

The probability distribution function of conditional quantiles will be denoted with $\hat{F}_{Y|X}$. It is obtained by taking the integral in (5), and substituting $F_Y^{-1}(\theta)$ with the expression for the estimator of conditional quantiles $Q_y(\theta)$, which were defined by equation (4)

(6) $$\hat{F}_{Y|X}\left(q_y \middle| X = x\right) = \int_0^1 \mathrm{I}\left[x\hat{\beta}(\theta) \leq q_y\right] d\theta = \sum_{j=1}^{J} (\theta_j - \theta_{j-1}) \mathrm{I}\left[x\hat{\beta}(\theta) \leq q_y\right]$$

13 We have to change the variable of integration, so that the random element will become the length of interval over which we integrate.

The shift from integration to summation is possible in a finite sample, since, as we have noticed earlier, there is a finite number J of distinct conditional quantiles, which are characterized by the J-vector of distinct solutions $\widehat{\beta}$. For a given x, we can consider (6) to be the conditional distribution of income implied by quantile regression. It follows from their definition that conditional quantiles can be estimated from the distribution of y given x as

$$(7) \qquad \widehat{Q}_Y(\theta|x) = \inf\{q_y : \textstyle\sum_{j=1}^{J}(\theta_j - \theta_{j-1})I[x\widehat{\beta}(\theta) \leq q_y] \geq \theta\}$$

Equation (7) is a convenient expression for conditional quantiles, since it is derived from their estimated conditional probability distribution. The next step is to obtain the marginal distribution of income, using some basic properties of probability. Denote the marginal density of the variable y by $f_Y(z)$. Notice that such density can be written as a function of the conditional density $f_{Y|X}(z|w)$, and of the covariates' density $f_X(w)$ as follows

$$f_Y(z) = \int_{-\infty}^{\infty} f(z,w)dw = \int_{-\infty}^{\infty} f_{Y|X}(z|w)f_X(w)dw$$

where $f(z,w)$ is used to denote the joint probability. In our framework, the conditional density of income can be derived from (6). By integrating such function with respect to

F_X, the result will be the unconditional distribution of y, denoted by $\hat{F}_Y$[14]

(8) $$\hat{F}_Y = \int \hat{F}_{Y|X}(q_y|X)dF_X = \int \left(\int_0^1 \mathrm{I}[x\hat{\beta}(\theta) \leq q_y]d\theta\right) dF_X$$

We now introduce some more notation. Let $i \in \{m, f\}$ represent male and female observations, so that we have two samples $\{(y_{ki}, x_{ki}): k = 1, \dots, n_i\}$ and the vector $\hat{\beta}$ can be estimated separately for the two groups. Moreover, consider x_{ki} as the kth row of the $n_i \times p$ matrix x_i. The distribution F_{X_i} of group i covariates can be approximated by the empirical distribution function as follows[15]

$$\hat{F}_{X_i}(z) = n_i^{-1} \sum_{k=1}^{n_i} I(x_{ki} \leq z), \qquad i \in \{m, f\}$$

where z is a given p-vector. Using the above distribution to evaluate F_{X_i}, we can estimate the unconditional distribution, expressing (8) as

(9) $$\hat{F}_{Y_i} = \int \hat{F}_{Y_i|X_i}(q_{yi}|X_i)dF_{X_i} = n_i^{-1} \sum_{k=1}^{n_i} \sum_{j=1}^{J} (\theta_j - \theta_{j-1})\mathrm{I}\left[x_{ki}\hat{\beta}_i(\theta) \leq q_{yi}\right]$$

In equation (9), each element of $\hat{\beta}_i$ is weighted by the length of the interval over which it

[14] We are exploiting the definition of marginal (unconditional) distribution, which is given by
$$F_Y = \int_{-\infty}^{\infty} \left(\int_{-\infty}^{\infty} f_{Y|X}(z|w) f_X(w) dw\right) dz = \int_{-\infty}^{\infty} \left(\int_{-\infty}^{\infty} f_{Y|X}(z|w)\, dz\right) f_X(w)\, dw = E_X\left[\int_{-\infty}^{\infty} f_{Y|X}(z|w)\, dz\right]$$

[15] Machado and Mata (2005) propose a random sampling method to approximate the covariates distribution. The two approaches are compared by Melly (2006), who shows that they become identical as the number of covariates' drows in the Machado-Mata procedure tends to infinity.

prevails. Moreover, each row of x_i contributes to the cumulative probability only if the resulting conditional quantile is lower than q_{yi}.[16] Unconditional (marginal) quantiles can now be estimated as

$$(10)\qquad \hat{y}_i(\theta) = F_{Y_i}^{-1}(\theta) = inf\left\{q_{yi}: n_i^{-1}\sum_{k=1}^{n_i}\sum_{j=1}^{J}(\theta_j - \theta_{j-1})\mathrm{I}\left[x_{ki}\hat{\beta}_i(\theta) \le q_{yi}\right] \ge \theta\right\}$$

Using these results, we can turn to the problem of decomposing group-wage differences. Assume that in a non-discriminatory labour market, females would have the same income's conditional distribution of males, that is, their characteristics would be rewarded as if they were males. The next step is to build the counterfactual quantile, in order to have a distributional measure of what would be female income, had the wage structure been the same as the male one. Such non-discriminatory quantile of income can be estimated by integrating $\hat{F}_{Y_m|X_m}$, that is, the male conditional income distribution, with respect to the female distribution of characteristics, $\hat{F}_{X_f}$. Using this procedure, we can estimate the following counterfactual income distribution, denoted by $\hat{F}_C$

$$(11)\qquad \hat{F}_C(q_c) = \int \hat{F}_{Y_m|X_m}\left(q_{ym}\middle|X_m\right)dF_{X_f} = n_f^{-1}\sum_{k=1}^{n_f}\sum_{j=1}^{J}(\theta_j - \theta_{j-1})\mathrm{I}\left[x_{kf}\hat{\beta}_m(\theta) \le q_C\right]$$

[16]Notice that the probability distribution (9) is not well defined if quantile-crossing occurs.

Equation (11) represents the non-discriminatory distribution of female income, assuming that the pay schedule faced by males would prevail in a fair labour market. Using the analogous procedure of equation (10), from the counterfactual distribution (11) we can obtain an estimator of the θth marginal quantile as follows

$$(12) \qquad \hat{y}_C(\theta) = inf\left\{q_C: n_f^{-1} \sum_{k=1}^{n_f} \sum_{j=1}^{J} (\theta_j - \theta_{j-1}) \mathrm{I}\left[x_{kf}\widehat{\beta}_m(\theta) \leq q_C\right] \geq \theta\right\}$$

Using the estimators in equations (10) and (12), we can carry out a wage gap decomposition similar to the traditional Oaxaca-Blinder decomposition, which was defined in equation (2). However, using the quantile regression approach, wage differences can now be evaluated at any θth quantile of the income distribution. More precisely, the pay gap between males and females can be divided in two parts, one representing the effect of different characteristics between the two groups, the other representing differences unexplained by the quantile regression model. For a given θ we can estimate

$$(13) \qquad \hat{y}_m(\theta) - \hat{y}_f(\theta) = [\hat{y}_m(\theta) - \hat{y}_C(\theta)] + [\hat{y}_C(\theta) - \hat{y}_f(\theta)]$$

It is useful to remark that the first addend is the so-called characteristics effect, since it is the consequence of the different distribution of covariates for the two groups. On the other hand, the second addend in (13) represents the so-called coefficient effect, since it is obtained by evaluating female characteristics using two different conditional distributions. The asymptotic distribution of the counterfactual estimator has been studied by Chernozhukov, Fernández-Val, and Melly (2013), who have shown the validity of exchangeable bootstrap inference procedures to estimate the covariance matrix.[17]

4 Data

To analyse the Italian gender wage gap, we will use the 2005, 2006, 2008 and 2010 waves of the ISFOL Population, Labour, Unemployment Survey (PLUS). These data are collected through telephone interviews, which are conducted during the first quarter of the year. Since the first PLUS survey of 2005, each year includes a proportion of panel observations, some of which are present in all four waves.[18] The target population is composed of individuals between 15 and 64 years old,[19] and the total sample size for each year is reported in the top part of Table 1.

The questionnaire is composed of specific sections designed to collect information on

---

[17] Drowing $n \times r$ observations with replacement from the empirical distribution to compute r estimates, is an example of a valid exchangeable bootstrap inference procedure.

[18] For a detailed illustration of the features of this survey and of the sampling design, see Mandrone (2012), chapter 9.

[19] In 2005, there were 38, 827, 322 individuals aged 15-64 in Italy, while the same population was 39, 655, 921 in 2010.

the following sub-groups of the population: young individuals between 15 and 29 years old; women between 20 and 49 years old; elderly population between 50 and 64 years old; unemployed individuals; employed population. A rich set of information for each of these categories is included, ranging from family characteristics to individual skills and personal history. It is then possible, for properly specified subsets of the sample, to provide a detailed explanation of some individual decisions (such as the choices of working, having children, studying) and to investigate which social environment factors may have influenced them.

Table 1: ISFOL PLUS Samples Size and Composition

	Year	2005	2006	2008	2010
Total ISFOL PLUS	*Males*	16.292	16.825	15.277	17.817
	Females	24.094	20.688	18.653	20.858
	Total	40.386	37.513	33.930	38.675
Employee Sample	*Males*	4.101 51,13%	3.765 64,89%	3.399 51,16%	3.948 39,97%
	Females	3.340 48,68%	2.716 66,27%	3.245 41,20%	2.923 39,51%
	Total	7.441 50,03%	6.481 65,47%	6.644 46,30%	6.871 39,78%

Notes: Criteria of selection of employee sample: full-time (30-72 h/week), wage at least 2 Euro/h, wage not imputed. Percentages refer to panel observations (i.e. employees whose wages are reported in more than one ISFOL PLUS wave).

The bottom part of Table 1 reports the composition of the sample that we will be studying. Despite the fact that also self-employed and those with project-linked job positions are present in the PLUS data sets, for our analysis we have considered only salaried employees, which form the largest category of workers. This choice is motivated by the fact that data on income is not harmonized among the above categories of job position. Moreover, in the context of gender wage gap estimation, we want to study a sample of workers that should be as homogeneous as possible. Indeed, the lack of overlapping support, an assumption discussed in Section 2, may become more relevant as we increase the categories of job position included. In order to gain further accuracy in our results, we have decided to include only full-time workers, and we have excluded all outliers and those who have chosen to not report their wages.[20]

We have used log-hourly net income (adjusted to the 2010 level) as the dependent variable. In Figure 1 we have graphed the Kernel density estimates of the wage distribution by year and by gender.[21] We can see that the modal observation is always lower for women, and that their income distribution tends to be shifted toward the left with respect to the male one. Both tendencies are a preliminary evidence of the presence of a gender pay gap. Such inpression is confirmed when looking at Table 2, where we can see that the average level of log hourly net wages is always lower for women. Notice also that 2008 seems to be the year with the lowest gender pay gap.

---

[20] More precisely, we have kept only individuals who were working between 30 and 72 hours a week, and whose net hourly wage was at least 2 Euro.

[21] All figures are placed in the Appendix.

Table 2: Mean and St. Dev. For Selected Variables

Group	Statistic	Log Hourly Wage	Age	Schooling	Long-Term Contract	Public Sector
			2005			
Males	Mean	2,108	41,532	12,473	0,894	0,367
	St. Dev.	0,373	12,728	3,369	0,308	0,482
Females	Mean	1,968	38,628	13,336	0,854	0,43
	St. Dev.	0,336	11,623	3,239	0,354	0,495
Total	Mean	2,045	40,228	12,86	0,876	0,396
	St. Dev.	0,363	12,328	3,339	0,33	0,489
			2006			
Males	Mean	2,096	40,991	12,578	0,881	0,355
	St. Dev.	0,371	12,893	3,259	0,324	0,478
Females	Mean	1,959	37,52	13,413	0,811	0,399
	St. Dev.	0,331	11,704	3,171	0,391	0,49
Total	Mean	2,038	39,536	12,928	0,852	0,373
	St. Dev.	0,361	12,525	3,248	0,355	0,484
			2008			
Males	Mean	2,062	41,409	12,766	0,852	0,326
	St. Dev.	0,377	13,259	3,275	0,355	0,469
Females	Mean	2,013	39,953	13,442	0,828	0,389

	St. Dev.	0,373	12,678	3,181	0,377	0,488
Total	Mean	2,038	40,698	13,096	0,84	0,356
	St. Dev.	0,376	12,998	3,247	0,366	0,479
			2010			
Males	Mean	2,079	41,872	12,751	0,848	0,315
	St. Dev.	0,385	13,524	3,244	0,359	0,464
Females	Mean	1,964	38,688	13,638	0,805	0,384
	St. Dev.	0,348	12,057	3,185	0,396	0,486
Total	Mean	2,03	40,517	13,129	0,83	0,344
	St. Dev.	0,374	13,016	3,249	0,376	0,475

Table 2 reports some descriptive statistics, by year and by gender, for the most important controls of our model. Notice that, in all years, women tend to be younger and better educated than men. Moreover, the proportion of long-term contracts is lower for the female group, while the proportion of public employees is lower in the male group. Notice also that mean log wages are always lower for female workers, while the dispersion of wages, as measured by the standard deviation, is similar for the two groups. Finally, there are no major differences in the sample composition across years, at least with respect to these variables.

As mentioned in Section 2, to carry out the decomposition exercise, we need to build a valid model of wage prediction. For this purpose, we have selected a rich group of independent variables. Specifically, the controls of the model are: years of schooling; a quadratic term for market experience, as approximated by age; a dummy for tenure, denoting employees who have been in their current job position for less than two years;[22] type of contract (long- or short-term); family characteristics (marital status, presence of pre-schooling age children, education of the mother); sector (services/goods production and public/private sectors); four occupation dummies and a dummy for firms with more than 50 employees; geographic variables (denoting people living in urban areas, people living in the North, with an interaction for North-West, and those living in the South, with an interaction for insular regions). In Section 6 we will compare the results obtained with this model specification to the ones obtained by adding nine occupation dummies, nineteen sectoral dummies and the entire set of Italian Regions.

5 Results

Our analysis of the gender wage gap is based on the method of equation (13). We have estimated the following decomposition

$$\hat{y}_{m,t}(\theta) - \hat{y}_{f,t}(\theta) = \left[\hat{y}_{m,t}(\theta) - \hat{y}_{C,t}(\theta)\right] + \left[\hat{y}_{C,t}(\theta) - \hat{y}_{f,t}(\theta)\right] \quad (14)$$

[22] The coefficients associated to dummies controlling for other levels of tenure were always not significant. Moreover, there were no substantial gender differences in the average levels of tenure.

for t = 2005, 2006, 2008, 2010. In order to approximate the various conditional distributions of income, we have estimated the quantile regression coefficients at 300 randomly drown percentiles.[23] The wage gap has been computed at 19 distinct quantiles. More precisely, we have estimated the decomposition model every 5 percentiles, starting from the 5th quantile until the 95th quantile. Finally, standard errors have been computed using 200 bootstrap replications.

The results of the decompositions of equation (14) are reported in Figures 2 and 3. For each year, the left graph represents the estimated total earning gap, as measured by the difference between the male and female conditional distributions of income. Instead, the graph on the right, is a decomposition of the total gender wage gap in a part attributable to individual characteristics, and a part attributable to differences in the estimated quantile regression coefficients (the wage structure or discrimination effect).

Notice that the estimated difference between male and female earnings is always positive and significant. With the exception of 2008, its magnitude is around 10 percentage points almost everywhere. Moreover, it tends to increase toward the top of

23 Using a subset of quantile regression solutions reduces the risks of quantile crossing. See footnote 10 on this point.

the wage distribution, an evidence that is usually described as a *glass ceiling effect*.[24] This tendency seems to be stronger in 2005 and 2006, since, in such years, the wage gap is above 20 percentage points for the highest percentiles of the earning distribution. Notice however that the shape of the gender earnings gap distribution is similar across years. This implies that women are discriminated more in jobs where the wages are higher.

It is interesting to compare the magnitude of the estimates across years. Notice that the earning gap was reducing between 2006 and 2008, but it has been increasing between 2008 and 2010, after that an economic downturn phase had begun. For example, the median estimate of the gap was around 10% in 2006, it dropped to less than 5% in 2008, and it increased again to more than 9% in 2010. These strong variations indicate that the wage gap is influenced by the economic cycle, and that while, before 2008, there was some progress during a phase of small growth, discrimination has increased again since the beginning of the economic crisis.

For what concerns the decomposition exercise, it is quite evident that the role played by individual characteristics is seldom significant. Indeed, the *characteristics line* of Figures 2

[24] This terminology has been introduced by Albrecht, Bjorklund, and Vroman (2003). To test the hypothesis of a *glass ceiling*, we have performed several tests on the equality of the estimated coefficients of the total wage gap at different percentiles of the wage distribution. In doing so, we have used the fact that each estimator is normally distributed around zero. The test on the equality of the estimated total wage gap between the 90th and at the 50th quantiles, as well as the test on the equality between the 75th and the 50th quantiles, is rejected in all four years.

and 3 is always very close to zero. This evidence shows that, in Italy, there has been a substantial convergence in human capital accumulation between men and women, so that gender differences can't be attributed to the levels of observable market potential. Moreover, the shape of the total wage gap resembles that one of the wage structure effect in all four years. This evidence suggests that the amount of discrimination is significant among Italian employees, and it is the main driver of the gender earnings differential. As we have noticed in Section 2, a significant wage structure effect can't be interpreted simply as a pure employers' disutility parameter, and it should be better described as a composite residual effect played by several unobservable market characteristics (which include discrimination as well). We can't identify directly the sources of this discrimination effect. However, in this and in the next Section, we will rule out some expanations and highlight some interesting tendencies.

To gain some more knowledge on what lies behind the wage gap, we have estimated the same decompositions of equation (14), dividing the sample between older employees, defined as those above age 35, and younger employees. The results of this exercise are plotted in Figure 4. It is quite evident that the gender earnings differential increases with age. Indeed, the estimated total wage gap almost doubles for older workers. Moreover, the *glass ceiling effect* is more pronounced for older workers, an evidence suggesting that women have difficulties in reaching those highly remunerated job positions which require more experience. Notice also that the pay gap was very close to zero for those with less than 35 years old in 2008, but it increased again in 2010. Also the level of discrimination is higher with age, as can be seen from the fact that the dotted line of Figure 4 (denoting discrimination among older workers) lies always above the dash-dotted line (which indicates discrimination among younger workers). Again, such findings suggest that women pay additional penalities as the level of market experience increases.[25]

With the possible exception of 2008, the distance between the wage structure effects and the total gap (represented by the solid and by the dashed lines) is higher among younger workers. This result is driven by the fact that, in our sample, younger female employees have better characteristics than older female workers, when compared with male employees of the same age. Therefore, there seems to be a higher level of human capital accumulation among young women, which could partly explain why the wage gap is smaller at earlier stages of the career. However, we should notice also that the

[25] In a recent article, Goldin (2014) explains the increasing relation between seniority and the gender pay gap by looking at the cost of flexibility. When women are older, their responsabilities in the informal labour market increase, and they become less willing to supply many hours of work. According to the author, this demand for flexibility in hours worked is costly for the employer. From our data, we could see that women supply less hours then man on average, but we could not find evidence of a positive relation between hours worked and hourly wages, even if measurement error might be negatively biasing our estimates.

discrimination effect continues to play an important role for younger women, an evidence that can't be fully explained by our data.

From the results presented here, it is quite evident that the wage structure effect is explaining most of the pay gap. In Section 6 we are going to test some hypothesis and provide further evidences on the existance of a gener pay gap. In particular, we will show that the earnings gap is mostly driven by *within occupation* and *within sector* differences in compensations. Indeed, we will construct several models with more covariates, and we will show that the estimated discrimination levels do not change much. Finally, we will deal with the problem of correlation between unobservable individual heterogeneity and *time-varying* covariates, showing that this issue is not affecting our results. However, we should be aware that the exogeneity of gender with respect to unobservable abilities remains an untestable assumption.

6 Robustness Checks

In Section 2 we have stated that, in order to have a meaningful decomposition of the gender pay gap, the ignorability assumption should be satisfied. Unfortunately, such assumption can't be tested directly. For this reason, we have stressed the fact that our measure of discrimination is best interpreted as a composite residual effect, which can't be fully explained by our data. Neverthelss, we have performed some robustness checks, in order to gain some knowledge on the quality of the wage prediction model estimated in the previous Section. Moreover, we have tested wether our estimates of discrimination are affected when correlation between *time-varying* dependent variables and individual fixed effects is taken into account, by looking at the direction and size of eventual biases.[26]

As a first robustness test, we have repeated the estimations carried out in Section 5, this time adding more explanatory variables to our model. In addition to the usual set of independent variables (which amounts to 26 covariates), we have included 9 dummies for occupational position, around 20 dummies for each Italian administrative Region, 17 dummies for different categories of economic sectors, and some more variables on the family background (education of the father). The results associated to this model, which contains a total of around 65 covariates, are represented by the dashed lines in Figures 5 and 6. The most important feature of this specification is that it allows to measure linearly any wage difference between a very detailed set of occupations and sectors. Therefore, our estimates become more suitable to take into account the effect segregation of women into less remunerative job positions. Finally, we have estimated

[26] This exercise is prone to type two errors. That is, whenever the more robust estimator (the Hausman-Taylor model in our case) is biased, no knowledge can be gained on the error associated to the less consistent estimator (the random effect model).

the quantile regression decomposition over a model where each of the 9 occupational dummies, and each of the 17 dummies denoting sectors were interacted. This model contains a total of around 220 covariates, and is represented by the dotted lines in Figures 5 and 6. This last specification allows us to capture also any non-linearity in the way occupational positions are rewarded across sectors.

In Figure 5, we have plotted the estimated gender wage differences that are attributable to differences in observable characteristics across groups. Each line represents a different model specification. In general, adding more covariates increases the amount of discrimination that is explained by characteristics. However, the difference between the 26-covariates and the other models are quite modest in magnitudes. The greatest difference is observable in 2005, where, for higher quantiles of the wage distribution, it reaches 5 percentage points. We should also stress the fact that this part of the income distribution is associated with the highest levels of estimated discrimination, and that this peak now seems to be in part driven by the composition of the sample in terms of occupational positions. For all other years, the differences in estimations across models seldomly exceed 2 percentage points.

In Figure 6, we have reported the estimated level of discrimination across model specifications. As can be noticed, the differences between the 26-covariates and the other models are smaller than in the previous case. In general, the wage structure effect is reduced when more covariates are added to the model. However, such characteristics effect lies always around or above 10%, with the only exception of 2008. Thus, we can conclude that most of the gender differences in pay originate *within* sectors and *within* occupations. With the partial execption of 2005, only a small percentage of the gender pay gap can be attributed to segregation effects, while discrimination continues to determine most of the differences in wages. In summary, the results presented in the previous section are quite robust when more detailed wage prediction models are considered.

As a final test on the validity of our model, we have exploited the longitudinal structure of the ISFOL PLUS data. We have tried to take into account individual characteristics, which are not observable in our data, but which may be correlated both with our explanatory variables and with wages. Indeed, a reason why the estimates of discrimination could be biased comes from the fact that some variables, such as schooling, experience or tenure, could be correlated with individual abilities not observable to the researcher. For example, using the ISFOL PLUS database, Borgna and Struffolino (2015) have shown that there are persistent gender differences in the dropout rates from secondary schooling. Thus there's the possibility that, due to the presence of unobservable dynamics that we fail to take into account, the same education level determines systematically different wage potential between women and men. Moreover, experience is approximated by age in our model, while actual experience could be influenced by time spent out of the working force. Alternatively, less productive workers could self-select into the category of workers with less years of tenure. All the above situations would lead to biased estimates of the coefficients associated to these variables, and could in turn undermine our decomposition exercise.

In general, correlation between individual fixed effects and *time varying* covariates is a quite common outcome,[27] and is confirmed in our data when performing a cluster robust Hausman test, which leads to a rejection of the null hypothesis. Unfortunately, the fixed effect model is not suitable for the estimation of the gender wage gap. Indeed, this model can't estimate the effect of *time invariant* characteristics on wages. In our context, the only feasible alternative is the Hausman-Taylor (HT) regression (Hausman and Taylor, 1981). This model allows for arbitrary correlation between *time varying* regressors and individual unobservable effects. For example, in this model returns to experience are measured as the average marginal effect on income of one more year of seniority, with respect to *individual-specific* wage means. However, the HT model also relies on a more demanding exogeneity assumption of the *time invariant* characteristics,

[27] See Polachek and Kim (1994) for a detailed discussion of the problem of unobservable individual heterogeneity in the context of the gender wage gap estimation.

which can't be tested in any way. Therefore, the effect of characteristics with limited, or no time variability at all, such as schooling and gender, are not consistently measured when individual fixed effects are correlated with them.

Table 4: Panel Estimates of Coefficients for Selected Variables

Model	**(1)**	**(2)**
Variable	*coeff./st. err.*	*coeff./st. err.*
Male	0.1195***	0.1200***
	(0.0043)	(0.0042)
Schooling	0.0201***	0.0265***
	(0.0007)	(0.0007)
Age	0.0183***	0.0267***
	(0.0015)	(0.0042)
Age Squared	-0.0001***	-0.0002***
	(0.0000)	(0.0000)
Observations	27,429	27,429

Notes: Model (1): GLS regression model, evaluated over the entire panel sample. Not consistent. Model (2): Hausman-Taylor regression, consistent to correlation between time-varying covariates and individual unobservable heterogeneity.

In order to see how much the estimated amount of discrimination is influenced by correlation of the *time varying* dependent variables with individual unobservable heterogeneity, we have compared the results of the HT model to the ones obtained from the inconsistent GLS estimator. The GLS estimator (also called random effects) is comparable to the OLS regression, but it is more efficient when dealing with a panel sample. This estimator relies on the assumption of exogeneity of all variables with respect to individual unobservable heterogeneity, an assumption that is not met in our data. Then, to obtain a consistent estimator for the coefficients associated to *time varying* variables, we have estimated the HT model. That is, we have first applied fixed effects regression, which is a consistent estimator in this context. Then, we have regressed the *time invariant* variables using the residuals of the first-stage fixed effects regression, employing the efficient GLS estimator.

In performing the longitudinal estimations described above, we have used our 26-covariates model, and we have pooled all four available years, in order to construct a panel sample. Finally, we have included a dummy equal to one for male observations. The coefficient associated to this variable represents the marginal effect of gender on income, keeping all else constant. Therefore, it can be interpreted as a measure of discrimination similar to the *wage structure effect*. Table 3 compares the estimated gender effect, together with a few selected covariates. As can be noticed by looking at the estimated effect of age on wages, the difference between the coefficients associated to *time-varying regressors* is quite significant across the two models. This result is confirmed by the Hausman test, which we have conducted over all such *time-varying* regressors. Instead, notice that the difference in the estimated gender effect is negligilble. This result suggests that correlation between *time varying* dependent variable and individual fixed effect is not biasing our estimations of gender discrimination. However, we must also interpret this result with caution. The possibility that also the HT model is providing biased estimates can't be neither tested nor excluded. Moreover, as we have already stressed at the beginning of this Section, unfortunaltely the validity of our *ignorability assumption* can't be tested. Nevertheless, gender discrimination seems to be an evidence very robust and persistent across models.

7 Conclusions

In this paper we have measured the gender wage gap using four waves of the ISFOL PLUS sample. This database is interesting for several reasons. First, it contains detailed information on workers' history, and it has a panel structure, which allows us take into account unobserved individual heterogeneity. Moreover, the chosen waves, which range from 2005 to 2010, cover a period that is interesting to study, due to the economic downturn which begun during 2008.

For our analysis, we have considered full-time private and public employees. We have adopted a quantile regression based methodology, which has allowed us to show that the wage gap is stronger in higher remunerated job positions. By dividing the sample between younger and older employees, we have shown that this earning gap is also increasing with age, a fact that can be in part attributed to an improvement of the level of human capital among younger women, while in part is the result of discrimination levels that are increasing with seniority.

In general, we have shown that women earn around 10% less than men, despite having similar market potential. This percentage is increasing with wages, and it reaches the level of 20% at higher quantiles of the income distribution. This *glass ceiling effect* is also stronger among older workers, a fact suggesting that women struggle to reach especially those well-remunerated positions that become available as experience increases. Moreover, we have shown that the level of discrimination is influenced by the economic cycle. Indeed, the gender wage gap was reducing in 2008, but, in 2010, it has reached again levels similar to the 2005 ones. This suggets that the economic recession, which was particularly severe in Italy during 2009, has had a negative impact on discrimination, worsening the position of women in the labour market, in terms of wages earned, more than that of men.

To further test our results, we have carried out some robustness checks. We have shown that increasing the information on workers' occupation, sector and location has only a small effect on the wage gap estimates. When including a model with a full set of sector-occupation interactions, the estimated effect of individual characteristics on the gender wage gap changed by no more than 5%, with this percentage being almost evrywhere around 1-2%. Moreover, the impact of additional dependent variables on the estimated wage structure effect was even smaller. Therefore, pay differences seems to be more relevant *within* similar job positions and sectors. Moreover, we have used the Haumsan-Taylor regression model to control for endogeneity problems that could arise with *time-varying* dependent variables. Our results show that this kind of endogeneity problem does not seem to be affecting the estimated gender wage gap. However, we were not able to test directly the assumptions required by the Hausman-Taylor estimator, and endogeneity problems could still be quite relevant even under this model.

We can conclude that women earn less than men, even when their characteristics are similar. The fact that gender differences in pay could not be attributable to the observable characteristics of our sample leaves the question on the source of this *wage structure gap* open. Based on previous literature,[28] several answers can be suggested,

[28] See for example Goldin (2014), Buser, Niederle, and Oosterbeek (2014), Booth, Francesconi, and Frank (2003), Sulis (2011).

even if they can't be fully tested from our data. In general, there could be a sorting mechanism, which makes men more likely to work for firms who pay higher wages.[29] Moreover, there could be differences in the way men and women are promoted, so that similar workers are paid differently due to internal compensation schemes. Women could be less willing to bargain and less competitive than men. They could prefer more flexible working hours, which could represent a cost for several firms. Finally, employers could be discriminating against women because of some form of disutility in hiring them, or because they could be in a monopsonistic position allowing them to pay lower wages and extract higher rents from female workers.

[29] As noted in Section 2, the previous, even if not vast, literature on this topic has not yet found evidences supporting this hypothesis

References

ABOWD, J., F. KRAMARZ, AND D. MARGOLIS (1999): "High Wage Workers and High Wage Firms," Econometrica, 67(2), 251–333.

ALBRECHT, J., A. BJORKLUND, AND S. VROMAN (2003): "Is There a Glass Ceiling in Sweden?," Journal of Labor Economics, 21, 145–177.

ASHENFELTER, O.C., H. FARBER AND M. RANSOM (2010): "Labor Market Monopsony," Journal of Labor Economics, 28(2), 203-210.

BECKER, G. (1973): The Economics of Discrimination. University of Chicago Press, Chicago, second edn.

BLINDER, A. (1973): "Wage Discrimination: Reduced Form and Structural Estimates," Journal of Human Resources, 8(4), 436–455.

BOOTH, A., M. FRANCESCONI, AND J. FRANK (2003): "A Sticky Floors Model of Promotion, Pay, and Gender," European Economic Review, 47(2), 295–322.

BORGNA, C., AND E. STRUFFOLINO (2015): "Early School Leaving Dynamics in Italy: The Heterogeneity of Gender Effects.," ISFOL Research Paper.

BUCHINSKY, M. (1998): "The Dynamics of Changes in the Female Wage Distribution in the USA: A Quantile Regression Approach," Journal of Applied Econometrics, 13, 1–30.

BUSER, T., M. NIEDERLE, AND H. OOSTERBEEK (2014): "Gender, Competitiveness, and Career Choices," The Quarterly Journal of Economics, 129(3), 1409–1447.

CARD, D., A. R. CARDOSO, AND P. KLINE (2013): "Bargaining and the Gender Wage Gap: A Direct Assessment," Discussion paper, IZA Discussion Paper.

CHERNOZHUKOV, V., I. FERNÁNDEZ-VAL, AND B. MELLY (2013): "Inference on counterfactual distributions," Econometrica, 81(6), 2205–2268.

CHRISTOFIDES, L., A. POLYCARPOU, AND K. VRACHIMIS (2010) "The gender wage gaps,'sticky floors' and'glass ceilings' of the European Union," IZA Discussion Paper No. 5044.

DI TOMMASO, M. L., AND D. PIAZZALUNGA (2013): "It is not a bed of roses. Gender and ethnic pay gaps in Italy," Discussion paper, University of Turin.

FORTIN, N., T. LEMIEUX, AND S. FIRPO (2011): "Decomposition Methods in Economics," in Labor Economics, ed. by O. Ashenfelter, and D. Card, vol. 4A of Handbooks in Economics, chap. 1, pp. 1–97. North-Holland, Amsterdam.

GOLDIN, C. (2014): "A Grand Gender Convergence: Its Last Chapter," The American Economic Review, 104(4), 1091–1119.

HAUSMAN, J. A., AND W. E. TAYLOR (1981): "Panel Data and Unobservable Individual Effects," Econometrica: Journal of the Econometric Society, pp. 1377–1398.

HECKMAN, J. (1979): "Sample Selection Bias as a Specification Error," Econometrica, 47(1), 153–161.

KOENKER, R., AND G. BASSETT (1978): "Regression Quantiles," Econometrica, 46(1), 33–50.

MACHADO, J., AND J. MATA (2005): "Counterfactual Decomposition of Changes in Wage Distributions Using Quantile Regression," Journal of Applied Econometrics, 20(4), 445–465.

MANDRONE, E. (2012): Labour Economics: PLUS Empirical Studies. Ediguida Editore, collana Temi e Strumenti.

MATANO, A., AND P. NATICCHIONI (2013): "Rent sharing as a driver of the glass ceiling effect," Economics Letters, 118(1), 55–59.

MELLY, B. (2005): "Decomposition of Differences in Distribution Using Quantile Regression," Labour Economics, 12(4), 577–590.

MELLY, B. (2006): “Estimation of Counterfactual Distributions Using Quantile Regression,” Review of Labor Economics, 68, 543–572.

NOPO, H. (2008): “Matching as a Tool to Decompose Wage Gaps,” The Review of Economics and Statistics, 90(2), 290–299.

OAXACA, R. (1973): “Male-Female Wage Differentials in Urban Labor Markets,” International Economic Review, 14(3), 693–709.

OAXACA, R., AND M. RANSOM (1994): “Discrimination and Wage Decomposition,” Journal of Econometrics, 61, 5–21.

POLACHEK, S. W., AND M.-K. KIM (1994): “Panel Estimates of the Gender Earnings Gap: Individual-Specific Intercept and Individual-Specific Slope Models,” Journal of Econometrics, 61(1), 23–42.

RANSOM, M. AND R.L. OAXACA (2010): “New Market Power Models of Sex Differences in Pay,” Journal of Labor Economics, 28(2), 267-289.

SULIS, G. (2011): “What Can Monopsony Explain of the Gender Wage Differential in Italy?”, International Journal of Manpower, 32(4), 446-470.

Appendix: Figures

Figure 1: Kernel Density Estimates of Hourly Wages by Year and Gender

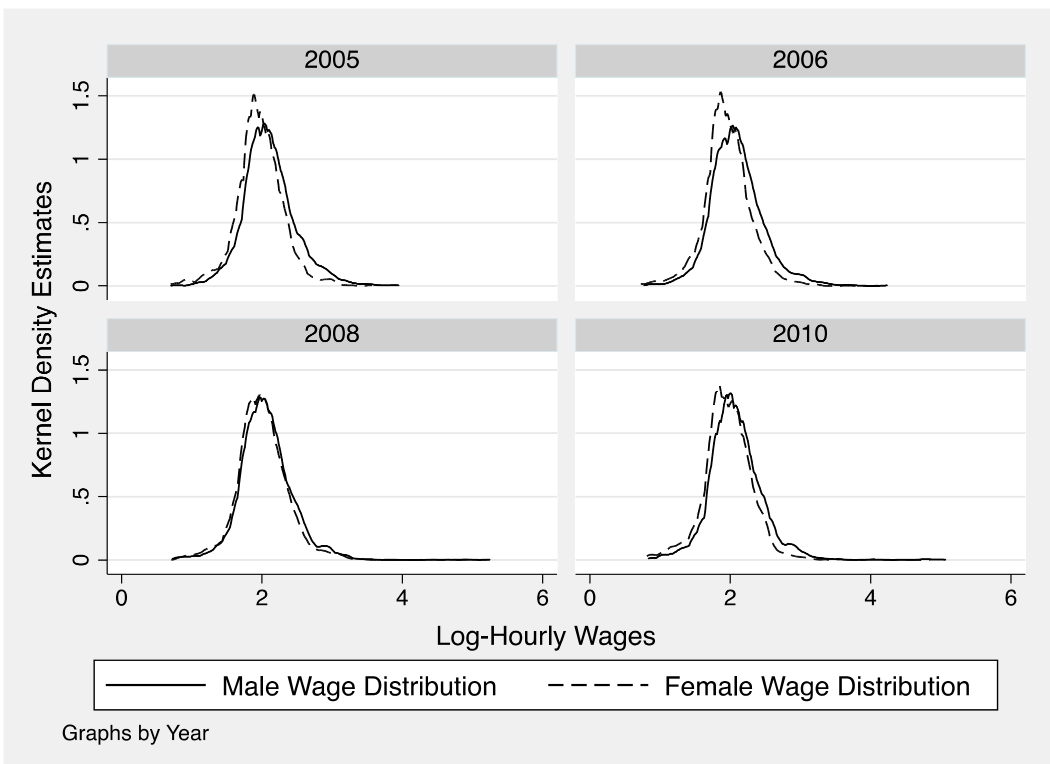

Figure 2: Quantile Regression Decomposition of the Gender Wage Gap by Year (1)

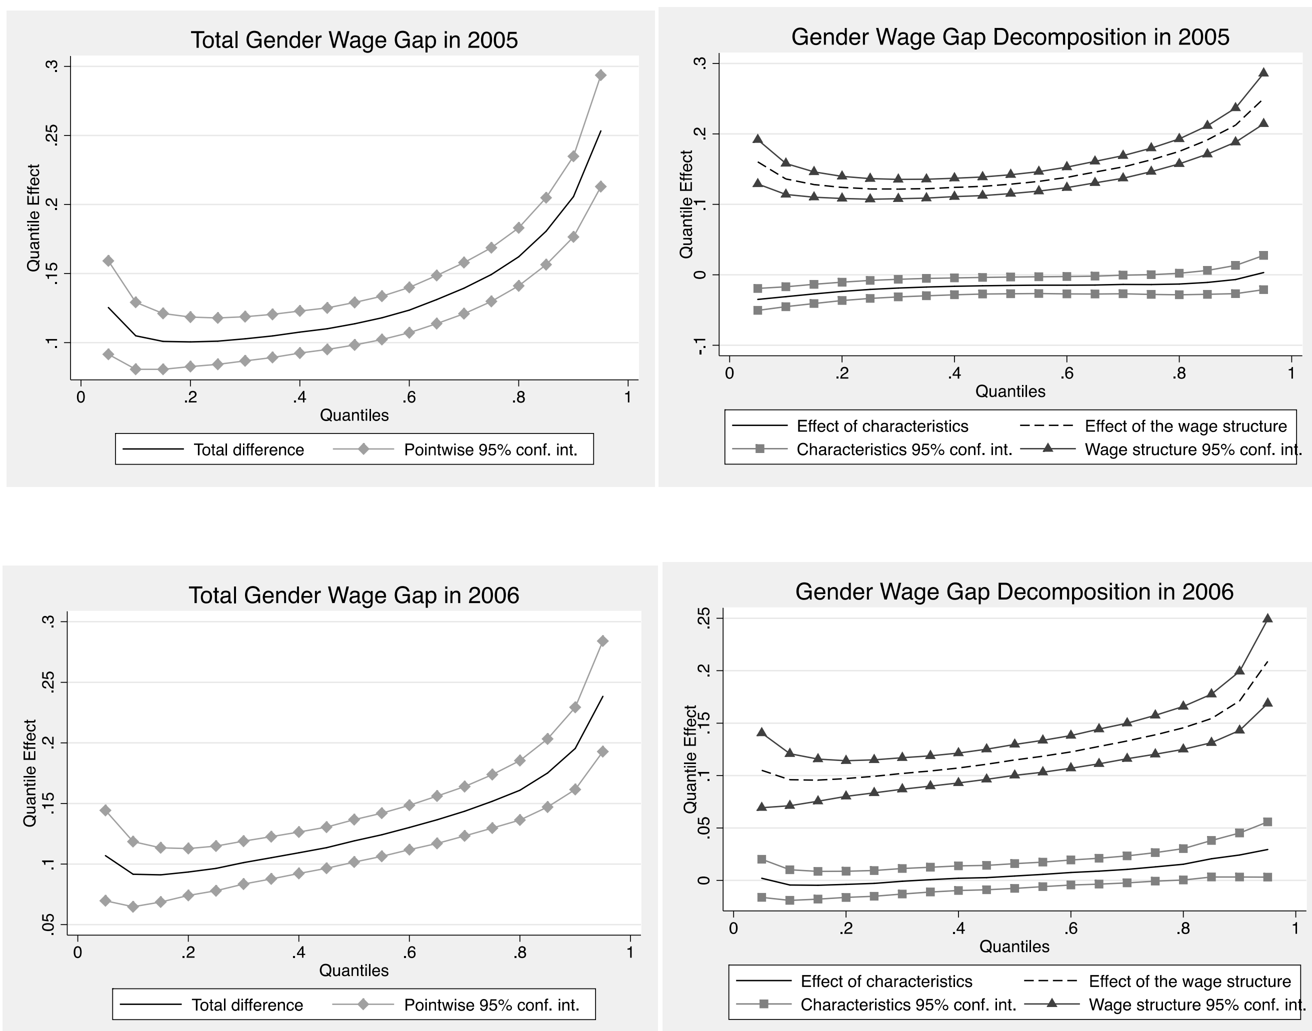

Figure 3: Quantile Regression Decomposition of the Gender Wage Gap by Year (2)

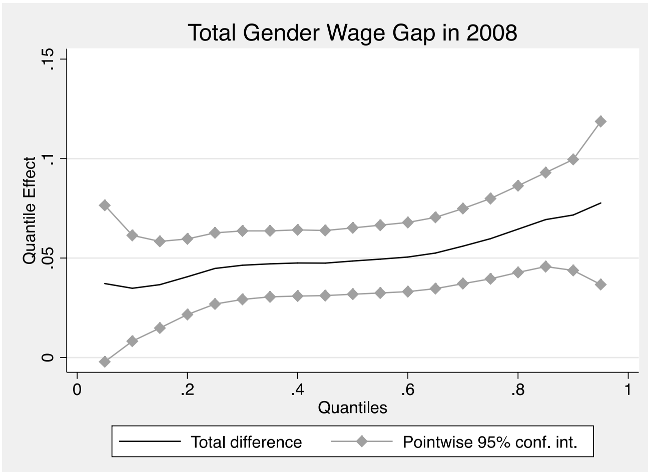

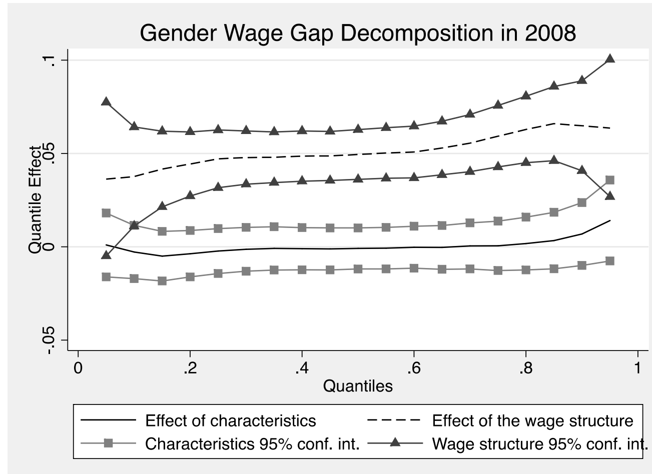

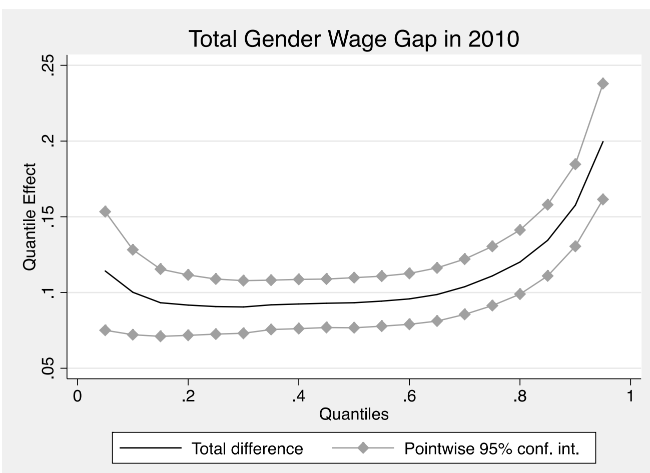

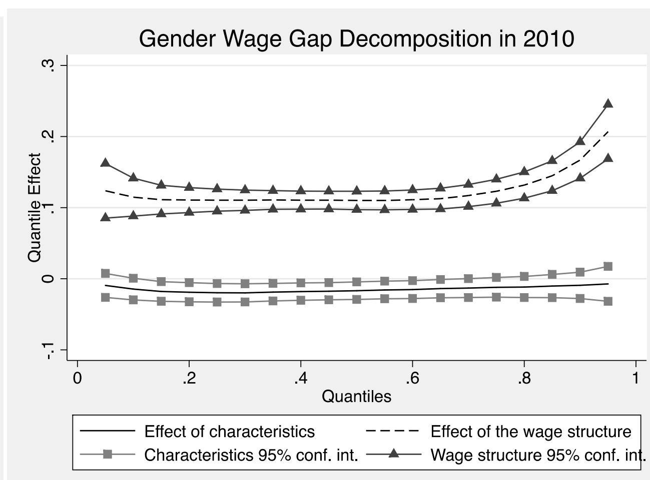

Figure 4: Quantile Regression Decomposition of the Gender Pay Gap by Age Group (Above and Below 35 Years Old)

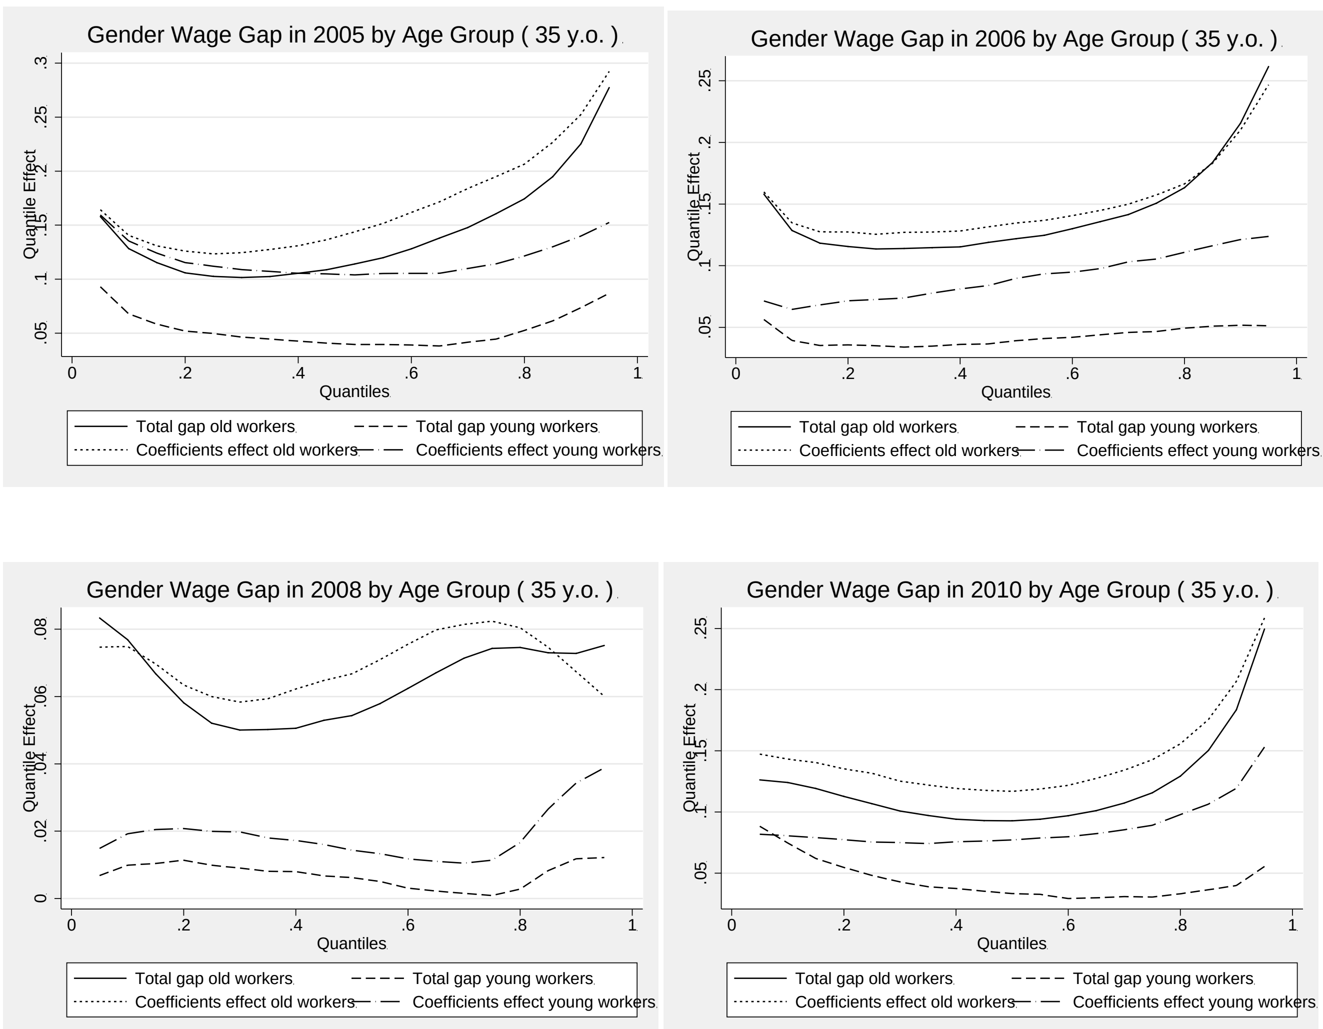

Figure 5: Characteristics effect obtained applying the counterfactual decomposition by year.

Each line reports results computed using a different model specification.

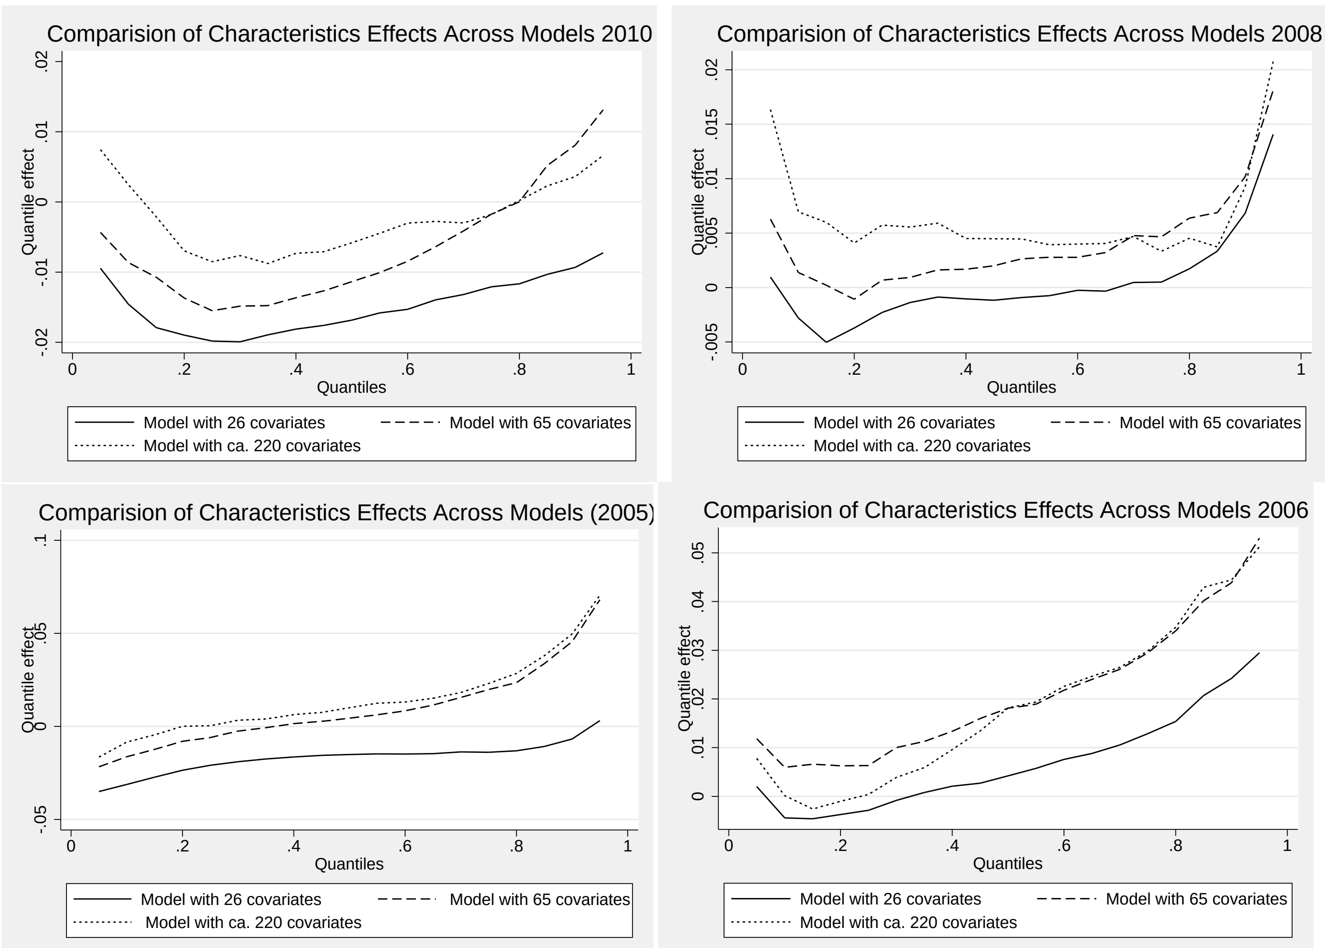

Figure 6: Wage structure effects obtained applying the counterfactual decomposition by year.

Each line reports results computed using a different model specification.